Contents

- Microsoft Message Analyzer Operating Guide
  - Getting Started with Message Analyzer
    - New Features and Updates
    - Installing and Upgrading Message Analyzer
    - Message Analyzer User Roles
    - Message Analyzer Feature Summary
    - Quick Session Startup
    - Technology Tutorials
      - Message Analyzer Tutorial
      - PEF Architecture Tutorial
      - ETW Framework Conceptual Tutorial
    - Message Analyzer Startup Options
    - Setting Message Analyzer Global Options
    - Accessibility
    - Procedures: Quick Start
  - Starting a Message Analyzer Session
    - Capturing Message Data
      - Targeting Live Data as an Input Source
        - PEF Message Providers
          - Microsoft-PEF-NDIS-PacketCapture Provider
          - Microsoft-PEF-WFP-MessageProvider
          - Microsoft-PEF-WebProxy Provider
        - Microsoft-Windows-NDIS-PacketCapture Provider
        - Understanding Event Parsing with a Provider Manifest
          - Generating a Provider Manifest
      - Configuring a Live Trace Session
        - Selecting a Trace Scenario
          - Built-In Trace Scenarios
        - Using a Custom Trace Scenario Template

Creating and Managing Custom Trace Scenarios
Adding a System ETW Provider
Modifying Default Provider Settings
Common Provider Configuration Settings Summary
Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog
Using the Advanced Settings- Microsoft-PEF-WFP-MessageProvider Dialog
System ETW Provider Event Keyword-Level Settings
Setting the Session Focus
Selecting Data to Capture
Specifying Advanced ETW Session Configuration Settings
Decrypting TLS and SSL Encrypted Data
Selecting a Session Data Viewer
Creating Remote Session Configurations
Configuring a Remote Capture
Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog
Performing a Live Capture
Procedures: Using the Network Tracing Features
Retrieving Message Data
Browse-Select-View Model
Targeting Saved Data as an Input Source
Configuring a Data Retrieval Session
Locating Supported Input Data File Types
Detecting and Supporting Message Truncation
Decrypting Input Data
Selecting Data to Retrieve
Applying a Session Filter to a Data Retrieval Session
Applying an Input Time Filter to a Data Retrieval Session
Specifying a Parsing Level
Selecting a Data Retrieval Session Viewer
Working With Special Input Requirements
Opening Text Log Files
Acquiring Data From Other Input Sources

Handling Azure Data
Loading System Event Log Data
Deriving Input Data with PowerShell Scripts
Loading SQL Data
Loading WPP-Generated Events
Loading OMS Log Data
Merging and Aggregating Message Data
Naming a Session
Performing Data Retrieval
Procedures: Using the Data Retrieval Features
Configuring Session Scenarios with Selected Data Sources
Editing Existing Sessions
Viewing Message Data
Data Viewer Concepts
Data Viewers
Analysis Grid Viewer
Saving Settings
Using the Field Chooser
Using and Managing Color Rules
Creating and Modifying Color Rules
Applying Color Rules
Managing Color Rules
Using the Find Message Feature
Using the Go To Message Feature
Applying and Managing Analysis Grid Viewer Layouts
Using the Analysis Grid Group Feature
Viewing OPN Source Code
Viewing Session Statistics and Progress
Grouping Viewer
Pattern Match Viewer
Using the Pattern Match Viewer
Understanding Message Pattern Matching

- Using the Pattern Editor
- Managing Pattern Expressions
- Gantt Viewer
- Chart Viewer Layouts
- HTTP Category
  - HTTP Content Type Payloads
  - HTTP Content Type Volumes
- General Category
  - Average Elapsed Time for Operations
  - Average Response Time for Operations
  - Cluster Levels
  - Event Log IDs
  - IP-Ethernet Conversations by Message Count
  - IP Ethernet Conversations by Message Count Top 20
  - TCP-UDP Conversations by Message Count
  - TCP-UDP Conversations by Message Count Top 20
  - Top Level Protocols Message Count
  - Top Level Protocols Message Count Over Time
- Network Category
  - IIS Log HTTP Traffic Volumes
  - IIS Log Server Bytes by Host over Time
  - IIS Log Top URI Bytes
  - IIS Log Top URIs by Time
  - TCP Rate and Diagnosis
  - TCP Stevens Graph
  - Top Talkers
  - Top Talkers Top 20
- Netlogon Category
  - Netlogon Message Types
- Networking Category
  - NTP Time Offset
- Common Category

Perfmon Log (.blg)
File Sharing Category
SMB File Stats
SMB Reads and Writes Bytes Sent
SMB Reads and Writes Bytes-Second
SMB-SMB2 Service Performance
SMB Top Commands
SMB Top Talkers
SysLog Levels
Interaction Viewer
Message Summary Tiles Viewer
Message Summary Lists Viewer
Perfmon Viewer
Charts (Deprecated)
Protocol Dashboard
Session Data Viewer Options
Common Data Viewer Features
Using the Filtering Toolbar
Applying and Managing Filters
Applying a Time Filter to Session Results
Applying and Managing Viewpoints
Working With Operations
Creating a Flat Message List
Filtering Data Sources
Setting Time Shifts
Configuring Time Format Settings
Using and Managing Message Analyzer Aliases
Creating Message Analyzer Aliases
Modifying Message Analyzer Aliases
Enabling and Disabling Message Analyzer Aliases
Refreshing Views
Performing Message Analyzer Operations with Aliases

- Managing Message Analyzer Aliases as Shared Items
- Configuring and Managing Message Analyzer Unions
  - Creating Unions
  - Modifying Unions
  - Refreshing Data Views Containing Unions
  - Performing Message Analyzer Operations with Unions
  - Managing Unions as Shared Items
- Viewing Process Name Data
- Tool Windows
  - Message-Specific Windows
    - Message Details Tool Window
    - Message Data Tool Window
    - Field Data Tool Window
    - Message Stack Tool Window
    - Selection Tool Window
    - Compare Fields Tool Window
  - Session-Specific Windows
    - Diagnostics Tool Window
    - Decryption Tool Window
  - Annotation Windows
    - Bookmarks Tool Window
    - Comments Tool Window
  - Other Windows
    - Field Chooser Tool Window
    - Output Tool Window
    - Session Explorer Tool Window
    - Map Tool Window
- Working with Message Analyzer Window Layouts
- Working With Message Analyzer Profiles
- Procedures: Using the Data Viewing Features

Analyzing Message Data

Filtering Message Data

- Filtering Loaded Input Data

- Filtering Captured Input Data
- Filtering Live Trace Session Results
- Filtering Column Data
- Writing Filter Expressions
  - Introduction to Creating and Applying Filters
    - Filter IntelliSense Service
  - Understanding the Filtering Language Basics
  - Using the Filtering Language
- Procedures: Using the Data Filtering Features

Saving Message Data

- Saving Session Data
- Selecting Messages to Save
- Saving Files in Native Format
  - Compatibility with Exported CAP Files
- Naming Saved Files

Automating Tracing Functions with PowerShell

Managing Message Analyzer Assets

- Sharing Infrastructure
  - Asset Manager
  - User Libraries
  - Managing User Libraries
- Managing Asset Collection Downloads and Updates
  - Syncing Items on First Startup
  - Filtering and Searching For Items
  - Downloading Assets and Auto-Syncing Updates
- Managing Microsoft OPN Parser Packages
- Managing the Default Subscriber Feed
- Creating Custom User Feeds
  - Manual Item Update Synchronization
- Sharing Asset Collections on a User File Share
- Procedures: Using the Asset Management Features

Extending Message Analyzer Data Viewing Capabilities

Configuring Chart Viewer Layouts
Using the Edit Chart Layout Dialog
Configuration Walkthrough of a Built-In Chart Viewer Layout
Managing Chart Viewer Layouts
Procedures: Using the Chart Configuration Features
Participating in the Message Analyzer Community
Message Analyzer Community Additions
Message Analyzer Team Blog
Message Analyzer Online Forum
Message Analyzer Feedback
Addendum 1: Configuration Requirements for Parsing CustomText Logs
Addendum 2: HTTP Status Codes

Microsoft Message Analyzer Operating Guide

17 minutes to read



Important
Microsoft Message Analyzer (MMA) is being retired and its download packages removed from microsoft.com sites on November 25 2019. There is currently no Microsoft replacement for Microsoft Message Analyzer in development at this time. If you already have Microsoft Message Analyzer installed, you may continue to use it, along with the OPN parsers you have already downloaded. Parsing ETW traces will also continue to work as before. After November 25 2019 when MMA is launched, it will attempt to connect to the back-end Feed service to check News and Assets updates and an error message will appear. To dismiss this errror message, see Dismiss Error Message.

Introduction
Microsoft Message Analyzer is a tool for capturing, displaying, and analyzing protocol messaging traffic, events, and other system or application messages in network troubleshooting and other diagnostic scenarios. Message Analyzer also enables you to load, aggregate, and analyze data from log and saved trace files. It is the successor to *Microsoft Network Monitor 3.4* and is a key component in the Protocol Engineering Framework (PEF) that was created by Microsoft to improve protocol design, development, implementation testing and verification, documentation, and support. With Message Analyzer, you can choose to capture local and remote traffic live or load archived message collections from multiple data sources simultaneously.

Message Analyzer enables you to display trace, log, and other message data in numerous data viewer formats, including a default tree-grid view, interactive Tool Windows, and other selectable graphical view **Layouts** that employ grids, bar element, timeline, and other visualizer components that provide high-level data summaries and other prominent statistics. You also have the option to configure your own custom **Layouts** for the **Chart** viewer. In addition, Message Analyzer now provides a **Profiles** feature, which creates interactive and integrated analysis environments that automatically display preset viewer and **Layout** configurations when data from specific input file types is loaded.

Quick Links
New Features and Updates — find out what's new in the latest release of Microsoft Message Analyzer.
Message Analyzer User Roles — determine your User Role and navigate to topics that help you get started with Message Analyzer in your role.
Message Analyzer Tutorial — take a detailed tour of Message Analyzer to learn about its capabilities, functions, and features.

Quick Session Startup — start a new Message Analyzer local trace session with a single click.
Procedures: Quick Start — see Message Analyzer in action right now by running several simple procedures.
Important for Network Monitor Users — review information about the differences between Network Monitor and Message Analyzer.
Feedback — provide feedback on any topic in this Operating Guide.

Other Links
Installing and Upgrading Message Analyzer — get a *free download* and install or upgrade Message Analyzer on your system.
Starting Message Analyzer for the First Time — learn about how to run Message Analyzer as an Administrator, security contexts and restrictions, and syncing Message Analyzer assets for automatic updates.
Training Videos — view several supplemental training videos to help you get started with Microsoft Message Analyzer.
Participating in the Message Analyzer Community — review options for participating in various Message Analyzer community venues.
Download PDF — download a PDF copy of this Operating Guide.

Information Roadmap

The topics outlined in this section provide a map into the documentation contained in the Message Analyzer Operating Guide. Use this map to quickly navigate to the topics that show you how to get started with Message Analyzer, how to use its basic and more advanced features, and to understand the underlying frameworks on which it is built. At a high level, the map breaks out into the three content spaces that are specified in the following table, within which you will find quick links that point to topics of interest in these spaces:

CONTENT SPACE	DESCRIPTION	NAVIGATION
Usage tasks	Review features and functions that you can use to perform various Message Analyzer operations.	Message Analyzer Usage Tasks
Usage procedures	Run procedures to see Message Analyzer in action and quickly familiarize yourself with its capabilities.	Message Analyzer Usage Procedures
Technology concepts	Review conceptual information to understand Message Analyzer features and the underlying technologies upon which they are built.	Message Analyzer Technology Concepts

Message Analyzer Usage Tasks

In this Operating Guide, Message Analyzer guidance is presented in the form of usage tasks. Each task provides some conceptual background with respect to the functions and features you will be working with, discusses how to use the associated UI features, and also includes example procedures to help you walk through various Message Analyzer usage contexts. To proceed directly to the usage tasks presented in this Operating Guide, click a task link below such as *Capturing Message Data*:

Getting Started with Message Analyzer

See the following topics to learn how to get started with Message Analyzer:

- **New Features and Updates** — read about the new and updated features in the latest release of Message Analyzer.
- Installing and Upgrading Message Analyzer — learn about Message Analyzer installation requirements, options, and other information, which includes upgrades from earlier Message Analyzer release versions, preserving user-created assets from prior installations, window docking layout changes, security contexts, and auto-syncing Message Analyzer assets for updates.
- **Message Analyzer User Roles** — determine your User Role based on how you intend to use Message Analyzer and navigate to topics that help you get started.
- Message Analyzer Feature Summary — review the main features of Message Analyzer and use the topic links to access more detailed feature descriptions.
- **Quick Session Startup** — learn about various methods you can use to very quickly start a new Message Analyzer session with a minimum of clicks.
- Technology Tutorials — read a tutorial on Message Analyzer functions before you dive into the usage tasks and procedures. Optionally, review the Protocol Engineering Framework (PEF) architecture and Event Tracing for Windows (ETW) framework tutorials to understand the technologies upon which Message Analyzer is built.
- Message Analyzer Startup Options — review the methods you can use to start Message Analyzer, which includes the arguments and command switches that are available to launch Message Analyzer from the command line.
- Setting Message Analyzer Global Options — set global options such as default values and settings that can affect Message Analyzer performance, display configurations, feature activations, and **Profile** activations.
- **Procedures: Quick Start** — run several simple procedures to quickly see Message Analyzer in action.

Capturing Message Data

Review the following topics to learn how to configure, start, and edit a Message Analyzer session, or configure a session scenario that targets multiple data sources, including local and multiple concurrent remote sessions. Discover how to start a session with a single click, how to use predefined **Trace Scenario** configurations and other message providers, how to create and save custom Live Trace Session configurations to run on-demand, how to use decryption, and how to enhance capture configurations with filtering, ETW system providers, and promiscuous mode:

- **Starting a Message Analyzer Session** — familiarize yourself with the types of sessions you can configure and start with Message Analyzer; also review common steps that you can use to create a basic session.
- Targeting Live Data as an Input Source — learn about the many different message providers that Message Analyzer uses as a source for live input data. Also learn about built-in **Trace Scenarios**, quickly starting a **Trace Scenario**, using system ETW providers, optimizing capture configurations, and session configuration workflow.
- **Configuring a Live Trace Session** — select and configure predefined **Trace Scenarios**, set predefined **Parsing Levels**, configure **Fast Filters** and **Session Filters**, configure system ETW providers, use advanced session configuration, select data viewers, and more.
- Built-In Trace Scenarios — review the functions and usage configurations of the built-in Message Analyzer **Trace Scenarios** in the **Network**, **Device**, **System**, and **File Sharing** categories.
- Decrypting TLS and SSL Encrypted Data — specify a server certificate and password to enable decryption and analysis of TLS/SSL encrypted traffic, which includes TCP, HTTP, and Remote Desktop Protocol (RDP) messages.
- Selecting Data to Capture — learn how to configure a Live Trace Session to capture specifically targeted data by applying a **Session Filter** and/or a **Parsing Level**.
- Configuring a Remote Capture — learn how to capture traffic concurrently on multiple remote hosts, which includes traffic on virtual machines that are serviced by a Hyper-V-Switch, along with advanced packet filtering and other special filters.
- Promiscuous Mode — learn how to capture data in P-Mode, if supported by your network adapter.
- Creating and Managing Custom Trace Scenarios — design a custom capture configuration template, save it as a **Trace Scenario**, and run it on demand.
- Editing Existing Sessions — learn how to reconfigure an existing session and apply the changes to existing data.
- Configuring Session Scenarios with Selected Data Sources — discover how to make use of the flexible session framework with multiple data sources capability that enables you to create Data Retrieval Sessions with multiple data loading configurations or Live Trace Sessions with multiple capture configurations for local and remote tracing.

Retrieving Message Data

Review the following topics to learn how to load input data from saved files, filter input data, and present it in a chosen viewer when loading messages through a Message Analyzer Data Retrieval Session:

- Browse-Select-View Model — learn about the Message Analyzer BSV infrastructure that enables you to browse for multiple data sources, filter or *select* specific data from those sources, and present results in a viewer of choice for data manipulation and analysis.
- Targeting Saved Data as an Input Source — browse for and load saved data from numerous log and trace file types into Message Analyzer, for example, *.matp, *.cap, *.evtx, .etl, *.log, *.csv, *.oms, *ps1, *.dmp, and *.saz files.
- **Configuring a Data Retrieval Session** — learn how to configure a Data Retrieval Session and make use of such features as session **Filtering**, **Truncated Parsing**, **Parsing Levels**, **Decryption**, Text log parsing, and more.
- Selecting Data to Retrieve — learn how to use a **Session Filter**, a **Time Filter**, and/or a **Parsing Level** to select specific data in a trace that you want to load into Message Analyzer.
- Acquiring Data From Other Input Sources — learn about other unique input data sources that Message Analyzer supports, such as Azure, Event Log, and SQL data; along with Operations Management Suite (OMS) log data and WPP-generated events. Also see Working With Special Input Requirements to learn about text log support.
- Selecting a Data Retrieval Session Viewer — learn how to specify a data viewer that displays message data that you load from one or more data sources in a Data Retrieval Session.
- Loading WPP-Generated Events — learn how to enable parsing of Windows software trace preprocessor (WPP)-generated events in Message Analyzer.

Viewing Message Data

Review the following topics to learn about the different data viewers that Message Analyzer provides, along with the capabilities that enable you to manipulate data views:

- Data Viewer Concepts — review background concepts about the Message Analyzer data viewing infrastructure to learn the basics on how data viewers work and interact.
- **Data Viewers** — learn about the data viewers that are available for data analysis, including the Analysis Grid Viewer, Grouping Viewer, Pattern Match Viewer, Gantt Viewer, and others, along with their associated data manipulation components. Also discover how to use Chart Viewer Layouts to display top-level protocol summary information and computed statistical values in graphic data visualizers for targeted analysis.
- Session Data Viewer Options — find out how to open various data viewers from multiple locations.
- Common Data Viewer Features — learn about Message Analyzer data manipulation tools that are common to the **Analysis Grid** and other viewers, such as view **Filters**, **Time Filters**, **Viewpoints**, **Operations**, **Aliases**, **Unions**, **Time Shifts**, and the **Flat Message List** feature that simulates the Network Monitor view.
- **Tool Windows** — understand how to use message-specific and session-specific **Tool Windows** that provide additional message details or configuration capabilities in Message Analyzer, for example, the **Diagnostics**, **Details**, **Message Stack**, **Decryption**, **Selection**, **Bookmarks**, and other **Tool Windows**.
- Working with Message Analyzer Window Layouts — learn how to create a customized working environment by selecting viewer and **Tool Window** preset configurations. Also learn how to use the redocking feature for data viewers and **Tool Windows**.
- **Working With Message Analyzer Profiles** — learn how to display a focused analysis environment by selecting a built-in or custom-designed data viewer and **Layout** preset configuration that automatically displays whenever you are loading data from a specific type of input file with which a **Profile** is associated.

Filtering Message Data

View the following topics to learn about selecting data in a Data Retrieval Session, applying filters to a Live Trace Session to isolate specific data, applying filters to trace results for analysis, using **Color Rules** to create conditional alerts or flags in a set of trace results, and understanding the Filtering Language:

- Filtering Loaded Input Data — apply a **Session Filter** to isolate specific data from a specified input file/s configuration.
- Filtering Captured Input Data — apply a **Fast Filter**, **Keyword** filter, **WFP Layer Set** filter, **Advanced Settings** filters, or an HTTP filter at the driver level to a Live Trace Session, or apply a predefined or custom Filter Expression as a **Session Filter** in the **New Session** dialog when configuring a Live Trace Session.
- **Filtering Live Trace Session Results** — select a filter expression from the common **Library** of predefined filters and apply it as a view **Filter** to the results of a Live Trace Session.
- Writing Filter Expressions — understand the Filtering Language so you can create your own filter expressions.

Saving Message Data

Review the following topics to learn how to save session data, which includes selecting messages to save, specifying the save file format, and using session naming conventions.

- Saving Session Data — read a quick overview of how to save your message data from a Data Retrieval Session or a Live Trace Session.
- Selecting Messages to Save — review the options that are available for saving message data.
- Naming Saved Files — review some naming strategies and other considerations for saving message data.

Automating Tracing Functions with PowerShell

Get a quick overview of the Message Analyzer functions that are enabled for the PowerShell scripting environment, as described in the following topics:

- Using PowerShell Cmdlets — read a synopsis of action, trigger, and other cmdlets that are available to automate various Message Analyzer functions and operations.
- Examining a PowerShell Script Example — review an example PowerShell script that configures a message provider, adds a **Trace Filter**, and sets various triggers for starting, filtering, stopping, and saving a trace session.
- Accessing PowerShell Cmdlets and Help — find out how to get PowerShell v3, access and update cmdlet help, and view the cmdlet help for Message Analyzer.

Managing Message Analyzer Assets

Review the following topics to learn about the Message Analyzer Sharing Infrastructure, user Libraries, automatic asset updates, downloading asset collections, and creating user feeds for sharing assets with others:

- Sharing Infrastructure — learn about the Message Analyzer Sharing Infrastructure; the user Library item collections that enable you to manipulate how data is captured, viewed, and analyzed; and how to manage these user Libraries.
- Managing Asset Collection Downloads and Updates — find out how to download user Library item collections and how to utilize the auto-sync feature to automatically receive user Library updates that are pushed out by a Microsoft web service.
- Managing Microsoft OPN Parser Packages — learn how to auto-sync updates to **OPN Parser** packages and download them from the Microsoft web service.
- Creating Custom User Feeds — create your own user feeds to which others may subscribe, for mutually sharing Message Analyzer assets with other team members, for example, **Filters**, **Trace Scenarios**, **Profiles**, viewer **Layouts**, and so on.
- Sharing Asset Collections on a User File Share — learn how to share user Library item collections directly with other users by exporting/importing collections or items to/from a file share.

Extending Message Analyzer Data Viewing Capabilities

Review the following topics to discover how to create custom **Layouts** for the **Chart** viewer that you can design to your own specifications with the use of various graphic visualizer components and data formulas. Enables you to extend Message Analyzer data viewing capabilities. Also learn how you can edit and customize any built-in **Layout** for the **Chart** viewer:

- Configuring Chart Viewer Layouts — learn how to use the Message Analyzer **Chart** viewer **Layout** configuration features to create a new **Layout** of your own design that is customized to your analysis environment.
- Using the Edit Chart Layout Dialog — learn how to use the controls and features of the **Edit Chart Layout** dialog to specify a visualizer component, data field values, and data formulas.
- Configuration Walkthrough of a Built-In Chart Viewer Layout — perform a walkthrough of the built-in **TCP/UDP Conversations by Message Count Layout** for the **Chart** viewer to familiarize yourself with the configuration features that you can use to create a functioning **Layout** of your own, based on a built-in and functioning **Layout**. Includes specifying graphic visualizer components and creating data formulas that perform various operations on message field values.

Message Analyzer Usage Procedures

If you want to proceed directly to usage procedures that demonstrate Message Analyzer features in the context of the usage tasks contained in this Operating Guide, click a link below:

Procedures: Quick Start — display saved data with the **Open** feature; start a Live Trace Session; display data quickly from your favorite **Trace Scenarios** by using the **Favorite Scenarios** feature on the Message Analyzer **File** menu or **Start Page**; load saved data through a Data Retrieval Session; and deploy various viewers, which includes **Layouts** for the **Chart** viewer, to display your data.

Procedures: Using the Network Tracing Features — run a **Local Network Interfaces** trace that isolates data to a particular network adapter and IPv4 address; perform a **Loopback and Unencrypted IPSEC** trace with a high-performance, driver-level **Fast Filter** that is set to capture HTTP traffic from TCP port 80; run a **Pre-Encryption for HTTPS** trace with driver-level **Hostname** and **Port** filters to isolate client and server HTTP message exchanges; capture traffic with a **Remote Network Interfaces** trace on a virtual machine (VM) that is serviced by a Hyper-V-Switch on a remote Windows 8.1, Windows 10, or Windows Server 2012 R2 host; and design a custom **Trace Scenario** and run it on demand.

Procedures: Using the Data Retrieval Features — browse for data and create a message collection to load into Message Analyzer; apply a **Session Filter** to loaded input data to isolate specific messages that you want to work with; display saved trace data in different viewers; use the **Recent Files** feature to display saved trace data to resume previous work; load data from multiple sources and save it as a single message collection; and apply a **Time Filter** to data being loaded into Message Analyzer.

Procedures: Using the Data Viewing Features — learn how to apply gradient-style **Color Rules** or a built-in view **Layout**; execute **Group** commands to group data and streamline message analysis; use the graphic visualizer components of the **Protocol Dashboard** to analyze top-level summary data such as top bandwidth consumption and message activity within a specified time window; analyze data with the interactive features of the **Protocol Dashboard** and **Analysis Grid** viewers; apply **Quick Filters** and **Viewpoints**; configure friendly **Aliases** for field values; create **Unions** of two or more message fields; and drive the display of various message details through **Analysis Grid** viewer and Tool Window interactions.

Procedures: Using the Data Filtering Features — create and apply filters to the data loading process, live captures, and trace results data to address and solve commonly encountered, real-world issues; create **Color Rules** to serve as an alert when certain message types, states, or values are present in a displayed message set, for example, TCP diagnostic information and SMB error status.

Procedures: Using the Asset Management Features — perform procedures that demonstrate how to manage user Library items and share them with others, or download and update Library item collections from the default **Message Analyzer** subscriber feed.

Configuration Walkthrough of a Built-In Chart Viewer Layout — perform a walkthrough of the built-in **TCP/UDP Conversations by Message Count** view **Layout** for **Charts** to familiarize yourself with the configuration features that you can use to create a functioning **Layout** of your own.

Message Analyzer Technology Concepts

If you want to expand your knowledge of the technologies upon which Message Analyzer is built, click the links below:

Technology Tutorials — get an overview of Message Analyzer functions and technology concepts, and learn about the PEF architecture and ETW framework components that support them:

Message Analyzer Tutorial
PEF Architecture Tutorial

ETW Framework Conceptual Tutorial

Getting Started with Message Analyzer

2 minutes to read

Message Analyzer takes new approaches to capturing, displaying, and analyzing message traffic, making it vastly different than other tools you may have used. Before you begin using Message Analyzer to capture live messages or retrieve data from saved message files and logs, you should familiarize yourself with its technologies and features. To advance your understanding of Message Analyzer and to get started quickly with its features, you are strongly advised to at least examine the feature summary and review the Message Analyzer Tutorial that are each described in this section. After doing so, you should give Message Analyzer a try by performing the *Quick Start Procedures* indicated below.

What You Will Learn

In the topics of this section, you will learn how to install and upgrade Message Analyzer, review a summary of the major features of Message Analyzer, review technology tutorials, learn about various options for starting Message Analyzer, and configure global Message Analyzer options, as indicated below.

In This Section

New Features and Updates — learn about the new features that are added to Message Analyzer since the last release, in addition to the latest updates to other product features.

Installing and Upgrading Message Analyzer — learn how to install and upgrade Message Analyzer, in addition to preserving user-created assets from an existing Message Analyzer installation.

Message Analyzer User Roles — define your user role based on the types of tasks you intend to perform with Message Analyzer and then navigate to topics that support your role.

Message Analyzer Feature Summary — review a feature summary and get an overview of Message Analyzer capabilities and functions.

Quick Session Startup — discover how you can very quickly acquire input data for Message Analyzer with as little as a single click.

Technology Tutorials — learn about Message Analyzer concepts, usage features, and the technologies on which they are built, for example, the underlying PEF architecture and ETW framework that support Message Analyzer operations.

Message Analyzer Startup Options — review various methods for launching Message Analyzer.

Setting Message Analyzer Global Options — learn about the global options that can affect Message Analyzer performance, display configurations, or feature activation.

Accessibility — review the accessibility shortcuts that are available with Message Analyzer.

Go To Procedures

To proceed directly to procedures that demonstrate some simple Message Analyzer tasks, see the **Procedures: Quick Start**.

New Features and Updates

4 minutes to read

This section provides a brief description of the ongoing new and updated features that are introduced in each release of Microsoft Message Analyzer.

New in Message Analyzer Version 1.4

The new and updated features that are provided in the Microsoft Message Analyzer version 1.4 release, are described in this section.

- **Filtering Toolbar** — consists of a new Filtering Toolbar that consolidates the major filtering functions of Message Analyzer into a common user interface (UI) that is accessible to all data viewers, including the **Grouping** viewer. Integrates the following filtering capabilities into the new toolbar:
  - View Filters
  - Time Filters
  - Viewpoints, including Disable Operations
  - Viewpoint Filters
  - Flat Message List (simulates the Network Monitor view)

  For more details about these feature, see the Using the Filtering Toolbar section.
- **Window Layouts** — organizes data viewers and **Tool Windows** into preset configurations for customizing your working analysis environment, with window layouts that range from simple to increasingly more complex configurations.

  For more details about this feature, see the Working with Message Analyzer Window Layouts topic.
- **Profiles** — enables you to utilize a set of built-in **Profiles** that contain specific data viewer and **Layout** presets that trigger whenever you are loading data from specific types of supported input files. Also enables you to create your own custom **Profiles**.

  For more details about this feature, see the Working With Message Analyzer Profiles topic.
- **Data Source Filter** dialog — originally part of the **Quick Filter** feature that was replaced by the **Time Filter** feature, which is now located on the Filtering toolbar that is accessible to every in-focus session viewer. You can display the **Data Source Filter** dialog by clicking the **Data Source Filter** item in the global Message Analyzer **Session** menu.

  For more details about this feature, see the Filtering Data Sources topic.
- **Chart configuration improvements** — provides a new UI for creating **Chart** viewer **Layouts** that simplifies configuration and editing. Also streamlines the use of visualizer components by limiting them to one per **Layout**.

  For more details about this feature, see the Extending Message Analyzer Data Viewing Capabilities topic.
- **Trace Session configuration improvements** — provides a new UI for adding system ETW Providers when configuring a Live Trace Session, along with a **Name** and **Guid** search capability that helps you find providers of interest. Also provided is the capability to specify GUIDs for custom providers that you can add to the **Available ETW Providers** list in the **New Session** dialog, while automatically registering them

with the operating system. Also relocates access to ETW session configuration on the **ETW Providers** toolbar in the **New Session** dialog.

For more details about these features, see Adding a System ETW Provider.

- **WPP setup simplified** — provides a simplified UI configuration in Message Analyzer **Options** to specify input TMF and PDB symbol files for parsing and viewing WPP-generated events.

  For more details about these features, see Loading WPP-Generated Events.

- **OMS support** — provides a new input data source that imports data from logs in an Operations Management Suite (OMS) Workspace for analysis.

  For more details about these features, see Loading OMS Log Data.

- **Pattern Match viewer improvements** — enables saving all **Pattern Match** messages as **Bookmarks**.

  For more details about this feature, see the Viewing Matched Instance Message Data and Bookmarks Tool Window topics.

- **Saving data in binary format** — enables saving trace data in binary format from the **Message Data Tool Window** by using the **Save Selected Bytes As** command.

  For more details about this feature, see the Message Data Tool Window topic.

- **Decoding URLs** — removes the special encoding characters in URLs and displays results in the **Decoded Value** section of the **Field Data Tool Window**.

  For more details about this feature, see the Decoding URLs topic.

- **Message time correlation improvements** — similar to **Time Offset** in Microsoft Network Monitor, a new **DeltaFromFirst** global property is provided for the **Analysis Grid** viewer to indicate the sequential running time of each message.

  For more details about this feature, see the **Tip** in the Default View Layout topic.

- **Display enhancements for data fields** — provides a new 8-bit field display for Flag type fields.

  For more details about this feature, see the Viewing Message Details Inline topic.

- **OPN viewer performance enhancements** — provides a more light-weight and better performing viewer for displaying OPN definitions.

  For more details about this feature, see the Viewing OPN Source Code topic.

- **Bookmarking improvements** — adds the capability to add bookmark messages to an existing **Bookmark**.

  For more details about this feature, see Adding Bookmarks to an Existing Bookmark or Group.

- **IP Address resolution** — Message Analyzer now resolves IP addresses to host names in the **Analysis Grid** viewer, in traces where this information exists.

  See the Analysis Grid Viewer topic for more information.

- **Promiscuous Mode selection enhancement** — provides simple selection of adapters that support P-Mode in the **Advanced Settings - Windows-NDIS-Packet-Capture** dialog, during Live Trace Session configuration, to enable capturing message data in Promiscuous Mode.

  For more details about this feature, see Capturing Data in P-Mode.

- **Enhancement to Pattern Expression configuration** — now supports timestamp calculations when creating new **Pattern Expressions** with the **Pattern Match Editor**.

For more information about creating Pattern Expressions, see Using the Pattern Editor.

- **Drag and Drop changes** — provides the capability to drag and drop multiple data files into Message Analyzer and aggregate all data into a single default data viewer.

  For more details about this feature, see Performing Data Retrieval.

- **Process Name detection improvement** — enables display of process name data natively in Message Analyzer for any ETW provider.

  For more details about this feature, see Viewing Process Name Data.

- **Assets** — provides additional assets that includes more **Layouts** for the **Analysis Grid** viewer, **Grouping** viewer, and **Chart** viewer. Also includes a simplified process for **Chart** viewer **Layout** configuration.

For more details about these features, see the following topics:

Applying and Managing Analysis Grid Viewer Layouts
Understanding the Built-In Grouping View Layouts
Chart Viewer Layouts
Extending Message Analyzer Data Viewing Capabilities

Installing and Upgrading Message Analyzer

11 minutes to read

IMPORTANT

The Microsoft Message Analyzer tool has been retired. We are leaving this page available for those who have downloaded the tool previously.

Installing Message Analyzer for the First Time

On this site, you will find details, new system requirements, brief installation instructions, and related resources such as **Known Issues** documentation and the Microsoft **Message Analyzer Operating Guide** in .pdf and .docx format. The Microsoft Message Analyzer release is available for installation in 32-bit and 64-bit versions. Installation requirements are listed in the table below for convenience.

Table 1. Message Analyzer Installation Requirements

COMPONENT	REQUIREMENT
Supported Operating Systems	**32-bit and 64-bit**: Windows 7, Windows 8, Windows 8.1, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, and Windows 10.
Redistributable Packages	**Minimum** — .NET Framework 4; **Recommended** — .NET Framework 4.5
Display Resolution	1024 x 768 or higher
Hard Disk Space	**Installation** — Minimum: 350 MB **Capturing and loading traces** — Recommended: 50 GB
RAM	**64-bit** — Minimum: 2 GB; Recommended: 8 GB **32-bit** — Minimum: 2 GB; Recommended: 4 GB
CPU	**Minimum** — 1.4 GHz, **Recommended** — 2 x 2.80 GHz (64-bit)

IMPORTANT

If you intend to perform long captures or load large traces, it is recommended that you use a 64-bit computer. Also, given that some Message Analyzer operations are CPU intensive, for example when filtering is combined with parsing, it is recommended that you use a dual-core processor for best performance.

Caution

If you are installing Message Analyzer on a Windows 7 computer, you might experience a random reset of the TCP stack and a TCP connection loss. If this impacts an application, you may have to restart it. Otherwise, there is a possibility that you might have to restart your computer.

> **NOTE**
> If you are installing Message Analyzer on a Windows 10 computer, the Windows Filtering Platform (WFP) message provider (Microsoft-PEF-WFP-MessageProvider) is installed as part of the operating system; therefore, the Message Analyzer installer (.msi) no longer installs this component. However, for all other supported down-level computers, such as Windows 8, the WFP message provider is installed with the Message Analyzer .msi application that you run from the download site. On Windows 10 computers, you can use this provider to conveniently capture traffic above the IP/Network Layer, which includes HTTP traffic, while removing lower-layer noise. Since the WFP provider is native to the Windows 10 operating system, you will not have to install any additional software to capture such traffic. To learn more about capturing this traffic on Windows 10 computers, see the WFP message provider PowerShell cmdlets on the TechNet Windows and Windows Server Automation with PowerShell site.

Upgrading an Existing Message Analyzer Installation
If you already have a Message Analyzer installation in place and you want to upgrade to the latest product version, read the **Warning** below and then proceed to Upgrading Message Analyzer.

> **WARNING**
> If you have Message Analyzer v1.0 (RTM) or later currently installed, it will be detected by the Message Analyzer installer application, in which case, you will be offered an **Upgrade** option. When you perform an upgrade, your existing Message Analyzer installation is removed, along with certain user assets you might have created, which includes any custom OPN parsers and OPN configuration files for text logs that are stored in the **OPNs** and **TextLogConfiguration\DevicesAndLogs** folders, respectively, under the following directory: `%LocalAppData%\Microsoft\MessageAnalyzer\OPNAndConfiguration\`. Moreover, if you happen to have a Message Analyzer beta version installed, you will need to manually remove it prior to upgrading Message Analyzer. When performing an uninstall in this case (from **Programs and Features** in **Control Panel**), you will also lose the aforementioned assets that you have created. Therefore, before you upgrade or manually uninstall a beta version of Message Analyzer, see Preserving User Created Assets Prior to Message Analyzer Installation to learn how to avoid the loss of user data.

Upgrading Message Analyzer

You may need to reinstall Message Analyzer, or you might simply want to upgrade to the latest Message Analyzer version to take advantage of recent changes and improvements. When you perform the Message Analyzer upgrade, you have the option to allow the Message Analyzer installer to detect your existing Message Analyzer v1.0 or later installation and offer the **Upgrade** option. However, you also have the option to perform a clean installation by manually uninstalling your existing version from **Programs and Features** in **Control Panel** before you run the Microsoft installer.. Also, you will definitely need to perform a clean installation if your existing Message Analyzer installation is a beta version, in which case the Microsoft Message Analyzer installer does not provide the **Upgrade** option. For this reason, you must be sure to manually uninstall your Message Analyzer beta version in the previously indicated manner prior to reinstalling.

> **IMPORTANT**
> Whether you manually uninstall Message Analyzer prior to reinstallation or if you perform an upgrade, ***all binaries and data are removed***, which might include some user-created assets, depending on where they are located, as described in Preserving User Created Assets Prior to Message Analyzer Installation

Preserving User Created Assets Prior to Message Analyzer Installation

Whether you need to manually uninstall Message Analyzer prior to reinstallation or if you will perform the **Upgrade** process, be sure to first back up any custom OPN parsers that exist in the default OPN parser directory indicated in the list that follows, as this user data will be lost. This applies equally to any OPN configuration files

for text logs that you have created and stored in their default location, also described in the list that follows. The only exception to this loss of data is if you have saved your custom OPN parsers and configuration files to a user-defined directory location, in which case they will not be deleted by a new installation of Message Analyzer. When your new Message Analyzer installation is complete, restore your custom OPN parser and OPN configuration files to the default directory locations so that Message Analyzer can locate and load them during application startup. The default directories are included in the list that follows:

- **Default OPN Parser directory** — the default location that Message Analyzer searches to load the default OPN parsers at startup is the following:
  `%LocalAppData%\Microsoft\MessageAnalyzer\OPNAndConfiguration\OPNs\` You can place all user-defined OPN parsers in this location and Message Analyzer will load them during startup; however, note that all parsers in this location will be overwritten if you manually uninstall/reinstall or **Upgrade** Message Analyzer.

  > **IMPORTANT**
  >
  > Please note that if you modify a default OPN parser in this directory, you will lose your changes when performing a new Message Analyzer installation. You will also lose your changes if an **OPN Parser** package that contains such a parser is set to auto-sync on the **Settings** tab of the **Start Page** in your current Message Analyzer installation and an update is pushed out by the sharing infrastructure. To avoid the possible loss of data in these circumstances, make a copy of the default OPN parser you wish to modify, rename it, and place it in a user-specified directory location, as described immediately below.

  > **NOTE**
  >
  > You also have the option to change the default OPN parser directory path that Message Analyzer uses, as specified in Message Analyzer Startup Options.

- **User-specified OPN Parser directory** — you can define an additional path that Message Analyzer will include when searching for OPN parsers to load during startup, as described in Message Analyzer Startup Options. In this case, you will need to start Message Analyzer from the command line and use the following switch to define the additional directory path: `/OPNLoadPathMerge=<path>` You can store all your custom OPN parsers in this location to prevent them from being deleted or overwritten.

- **Default OPN Configuration files for text logs directory** — the single, default location in which Message Analyzer searches for text log configuration files is the following:
  `%LocalAppData%\Microsoft\MessageAnalyzer\OPNAndConfiguration\\TextLogConfiguration\DevicesAndLogs\` You can place all custom OPN configuration files for parsing text logs that you create in this directory location; note, however, that all configuration files in this location will be overwritten by a new Message Analyzer installation. The configuration files will also be overwritten if you have the **Device and Log File Version 1.3** asset package set to auto-sync on the **Settings** tab of the Message Analyzer **Start Page** and an update is pushed out by the sharing infrastructure.

  Caution

  To avoid this loss of data, you are strongly advised to make a separate backup of any OPN configuration files for parsing text logs that you have created.

> **NOTE**
>
> All other *user-created* assets such as **Filters**, **Charts**, **Aliases**, **Color Rules**, **View Layouts**, and so on, are stored in the following directory: `%AppData%\Microsoft\MessageAnalyzer\My_Items\` This location is not impacted by Message Analyzer upgrades.

Message Analyzer Window Docking Layout

Message Analyzer now uses a versioned docking layout configuration file to keep track of the locations in which you dock various windows in the user interface (UI), for example, the **Tool Windows**. When you shut down the Message Analyzer application, the file is updated with information for the current window docking configuration. This tracking enables you to maintain a consistent window layout in the UI across every Message Analyzer restart in the same installation. The configuration file is located in the following directory:
`%LocalAppData%\Microsoft\MessageAnalyzer\app.WindowLayoutAsset.cfg`

When you perform an upgrade to the latest version of Message Analyzer, any existing docking layout configuration file in this location will be overwritten by a new default docking layout configuration file that comes with the your upgrade installation. As a result, you will lose your old window docking layout configuration in the Message Analyzer UI.

Installing Message Analyzer From the Command Line

You have the option to run the Message Analyzer installer (.msi) from the command line. To do this, you will need to save the .msi from the Microsoft Download site, rather than run it. After you save it, you can execute the following command at a command prompt to install Message Analyzer: `msiexec /i MessageAnalyzer.msi`

One reason that you might want to install Message Analyzer from the command line is that you can create an installation log, which can provide information about the install if you are having any issues:
`msiexec /i MessageAnalyzer.msi /l* <logfilePath>`

To review additional installation options, see the msiexec help by running the following command string at a command prompt: `msiexec /?`

Starting Message Analyzer for the First Time

When you are ready to start Message Analyzer for the first time, go to the **Start** menu, **Start** page, or task bar of your computer and click the **Microsoft Message Analyzer** icon to launch Message Analyzer. For example, on computers running the Windows 10 operating system, click the **Start** page and type the text "Message Analyzer" to display the **Message Analyzer** program icon. To ensure that you will have access to all the critical Message Analyzer features and capabilities, start Message Analyzer in administrative mode by right-clicking its icon and then selecting the **Run as administrator** option.

However, if you are starting Message Analyzer immediately after installing it, you should log off and back onto your computer first, and then start Message Analyzer as indicated. This action ensures that in all subsequent logons following installation, your security token will be updated with the required security credentials from the Message Capture Users Group (MCUG) so that you can capture network traffic in **Trace Scenarios** that use the **Microsoft-PEF-NDIS-PacketCapure** provider, **Microsoft-Windows-NDIS-PacketCapture** provider, and the **Microsoft-PEF-WFP-MessageProvider**.

Note also that when you start Message Analyzer for the first time, you will be prompted to opt in or out of automatic asset updates from a dedicated Microsoft web service, as described in Syncing Message Analyzer Assets for Automatic Updates below.

Security Contexts and Restrictions You should also be aware that when you run Message Analyzer in the administrative mode, it can result in varying security contexts between applications. For example, this means that you will be unable to use the Message Analyzer drag-and-drop feature to open saved trace and log files while running in administrative context.

IMPORTANT

Even if you log off your system, log back on, and receive the required security credentials from the MCUG, you will still need to use the **Run as administrator** option if you want to capture message traffic in Message Analyzer **Trace Scenarios** that use the **Microsoft-Windows-NDIS-PacketCapture** provider or the **Microsoft-PEF-WFP-MessageProvider**, which both have remote capabilities. Because of the inherent remote capabilities of these message providers, security restrictions must be applied.

Syncing Message Analyzer Assets for Automatic Updates The first time that you start up Message Analyzer, you have the option to configure the Automatic update of Message Analyzer assets, such as **Filters**, **Trace Scenarios**, **Viewpoints**, **Chart** viewer **Layouts**, and so on. If you opt in to automatic updates, no further action is required, as the update process occurs silently in the background whenever you start up Message Analyzer and asset updates are available. If you do not opt in, then the asset versions that install with the latest version of Message Analyzer will not be updated by the Microsoft web service that handles the installation of such asset versions. However, you can set any asset collection or all of them to the auto-sync state anytime thereafter with the use of the **Asset Manager** dialog.

More Information

To learn more about syncing Message Analyzer assets for automatic updates at first startup, see Syncing Items on First Startup.

To learn more about the **Asset Manager** dialog, see the Asset Manager topic.

See Also

Quick Session Startup

Message Analyzer User Roles

4 minutes to read

The Message Analyzer user audience primarily consists of network administrators, support analysts, and developers. However, other personnel who support this audience often use Message Analyzer in a limited capacity that is restricted to the task of taking traces and forwarding the data to others for analysis. The list that follows briefly describes the types of tasks that typical Message Analyzer users perform.

- **Network Administrator** — uses Message Analyzer as a network troubleshooting and analysis tool.
- **Network Support Analyst** — uses Message Analyzer in Help Desk scenarios to capture data for clients or to aggregate log and trace data from multiple sources and different time zones for analysis.
- **Protocol Developer** — generates network protocol code and uses Message Analyzer as a validator of protocol behavior, architecture, message field values, and state.
- **Event Tracing Developer** — instruments applications with ETW technology and uses Message Analyzer to capture the events from ETW providers.
- **Tracer** — typically a client who is experiencing issues in some area and is requested to take a trace for in-depth analysis by a Network Administrator or Support Analyst.

The diverse audience described above can be represented in two high-level categories of User Roles, basic and advanced, as described immediately below. The list that follows briefly summarizes each high-level User Role and provides pointers to the appropriate topics in this Operating Guide for each User Role to get started with Message Analyzer. The information in this list should help you determine the User Role that best fits your use of Message Analyzer.

- **Basic User Role** — if you are a **Tracer** and you intend to use Message Analyzer only to take traces, save them to file, and then pass them along to colleagues, administrators, or support personnel for analysis, proceed to the following topics to get started quickly with these tasks.
- Quick Session Startup — describes all the options for starting a trace session. For example, you can start a Live Trace Session immediately with no session configuration required, by clicking the **Start Local Trace** button on the Message Analyzer **Start Page** to capture messages on the local computer at the Link Layer. You can also start a Data Retrieval Session to import saved data into Message Analyzer, with no session configuration required, by clicking the **Open** button on the Message Analyzer **Start Page** and then navigating to saved files for selection.
- Working with Message Analyzer Window Layouts — optionally, create a basic data display environment by choosing one of the simple built-in **Window Layouts** that organizes data viewers and **Tool Windows** into a preset configuration.
- Performing a Live Capture — provides a summary of two methods that you can use to start a Live Trace Session quickly, which includes using the **Start Local Trace** and **Favorites** features.
- Performing Data Retrieval — describes additional options for retrieving input data for Message Analyzer, including loading data from text logs.
- Procedures: Quick Start — provides several examples of simple procedures that you can perform to see Message Analyzer in action, although some minor configuration is required in these scenarios.
- Saving Session Data — describes how to save Message Analyzer trace data, how session results are saved, and also specifies the default save location.

- **Advanced User Role** — if you are a **Network Administrator**, **Network Support Analyst**, or a **Developer**, and you intend to use Message Analyzer to create custom trace configurations, run live traces, retrieve saved logs and trace file data, and perform analysis, review the following topics to get started quickly with these tasks.
- Quick Session Startup — review this topic to familiarize yourself with various methods for starting a Message Analyzer session quickly.
- New Features and Updates — review this topic to obtain an overview of the new and updated features that exist in the current release of Message Analyzer, so that you can leverage them in your work.
- Starting a Message Analyzer Session — review this topic to obtain the following:
    - An overview of the configuration options that are available when you are creating a Live Trace Session or Data Retrieval Session.
    - An overview of the workflow for creating new Message Analyzer sessions that capture live data or retrieve saved data.
    - Links to topics that provide in-depth details about capturing and retrieving message data, configuring session scenarios, editing existing sessions, and running procedures that provide examples of how to use the network tracing and data retrieval features. Topics of particular interest might be the following:
        - Targeting Live Data as an Input Source
        - Configuring a Live Trace Session
        - Performing a Live Capture
        - Targeting Saved Data as an Input Source
        - Configuring a Data Retrieval Session
        - Performing Data Retrieval
- Viewing Message Data — review this topic to learn about a wide assortment of data viewing options that are available for live or saved data, including built-in data viewers, **Chart** viewer **Layouts**, **Tool Windows**, **Window Layouts**, **Profiles**, and other data manipulation tools. Also run procedures that provide examples of how to use data viewers and **Tool Windows**, and how to apply **Filters**, **Viewpoints**, and data **Grouping**. Topics of particular interest might be the following:
    - Data Viewers
    - Using the Filtering Toolbar
    - Tool Windows
    - Working With Message Analyzer Profiles
- Analyzing Message Data — review this topic for an overview of Message Analyzer analysis capabilities. Given that viewing message data and analyzing message data are closely related, this topic will give you a better idea of how to use the Message Analyzer viewing features as analysis tools.
- Saving Session Data — as described earlier.

> **NOTE**
>
> If your tasks fall into the **Advanced User Role** category, you are also encouraged to review the Message Analyzer Tutorial to obtain a broad understanding of Message Analyzer features and functions.

Message Analyzer Feature Summary

27 minutes to read

Microsoft Message Analyzer contains a broad and versatile range of features that build upon and exceed many of those of its predecessor, Microsoft Network Monitor. These features are designed to improve your usability experiences and to expand your capabilities set when loading, capturing, analyzing, and troubleshooting message traffic with Message Analyzer. The following provides a summary of these features that is broken down according to the means by which you can access them.

Global Menus

Message Analyzer provides global menus for quick access to various features that you will regularly use. The global menu names are located in the upper-left section of the Message Analyzer user interface and are described in the following sections:

File Menu The global Message Analyzer **File** menu provides access to the features described in the list that follows.

- **New Session** — click this item to open the **New Session** dialog, from where you can choose a source from which to acquire data; for example a **Live Trace** or saved **Files**. Clicking the **New Session** item also displays a submenu that contains various items that determine what type of session you will start. With the exception of the first two bullet items immediately below, the remaining submenu items also appear in the **New Session** dialog as **Data Source** buttons that you can click to begin the configuration of a new session, based on the type of input message data you want to acquire:
  - **Blank Session** — opens the **New Session** dialog from where you can select a data source under **Add Data Source** to use as input to Message Analyzer.
  - **From Current Session** — opens the **New Session** dialog to a new session configuration that derives configuration settings from the current in-focus session.
  - **Live Trace** — opens the **New Session** dialog with the **Live Trace** tab selected, from where you can specify one or more target computers on which to capture data; select a built-in **Trace Scenario** from the scenario Library; and configure various provider settings and filters to customize your trace configuration before starting a live trace. The **New Session** dialog also enables you to specify global session settings such as a **Session Filter**, **Start With** data viewer selection, and **Parsing Level**. You also have the capability to run multiple concurrent Live Trace Sessions with different message providers on different target computers by adding one or more **Live Trace** input sources by clicking the **New Data Source** tab and specifying the hosts from which to capture the data. You can also use a single session with a specified message provider to collect data from multiple specified host machines.

More Information
To learn more about starting a new Live Trace Session, see Starting a Message Analyzer Session.

> **NOTE**
> If you intend to *capture* messages that are encrypted with the Transport Layer Security (TLS) or Secure Sockets Layer (SSL) security rotocols, for example, HTTPS and Remote Desktop Protocol (RDP) messages, you have the option to enable any Live Trace Session for *Decryption** so that you can view the decrypted data along with decryption session statistics. For more information, see Decrypting TLS and SL Encrypted Data.

- **Files** — opens the **New Session** dialog with the **Files** tab selected, from where you can configure a Data Retrieval Session to acquire data that exists in one or more saved files. You can also select specific data to retrieve from such sources by using filters, for example a **Time Filter** and/or **Session Filter**.

  A **Truncated Parsing** check box is also included in the **Files** tab configuration to indicate when truncated messages exist in files from which you are retrieving data, at which time Message Analyzer switches to a pared-down truncation parser set. You have the option to unselect this check box or to select it manually if Message Analyzer did not automatically detect truncated messages.

  More Information
  To learn more about starting a new Data Retrieval Session, see [Starting a Message Analyzer Session] starting-a-message-analyzer-session.md).

  > **NOTE**
  > If you intend to *retrieve* messages that are encrypted with the Transport Layer Security (TLS) or Secure Sockets Layer (SSL) security protocols, for example, HTTPS and Remote Desktop Protocol (RDP) messages, you have the option to enable the Data Retrieval Session for **Decryption** so that you can view the decrypted data along with decryption session statistics. For more information, see Decrypting TLS and SSL Encrypted Data.

  In addition, the **Files** tab configuration provides you with the capability to retrieve data from textual log files and to select from a list of configuration files that support log file parsing. The **Truncated Parsing**, **Decryption**, and text log parsing features are described in Configuring a Data Retrieval Session.

- **Azure Table** — opens the **New Session** dialog to a configuration interface that enables you to specify an **Account Name**, **Account Key**, and **Table Name**, from which you can load Azure event log data into Message Analyzer.

- **Event Logs** — opens the **New Session** dialog to the **Event Logs** tab, which contains a large list of event logs that were generated on your computer. You can select one or more of the event log check boxes and click **Start** to retrieve the data from the selected logs.

- **PowerShell** — opens the **New Session** dialog to the **PowerShell** tab, from where you can write a PowerShell query. If your PowerShell script obtains data from a remote source, you can specify a **Host** name and connection credentials in the **New Session** dialog. Note that if your PowerShell script captures event or network traffic and logs to an event log (*.etl) or Message Analyzer *.matp file, respectively, you can import the data from such files through a subsequent Data Retrieval Session.

- **Sql** — opens the **New Session** dialog to the **Sql** tab, from where you can load data into Message Analyzer from any SQL database table. Provides facilities to specify a **Connection** string, user credentials, **Table** name, **Timestamp**, and a **WHERE** clause.

- **Oms** — opens the **New Session** dialog to the configuration interface for access to Operations Management Suite (OMS) logs. Specify user credentials to log on to Azure, select a **Subscription**, select a **Workspace** that you have configured, and then specify an Azure Resource Manager (ARM) query to extract OMS data from logs that are written by solutions that are enabled in a Workspace you created.

- **Open** — provides a submenu with the following two items:
  - **From File Explorer** — click this item to launch Windows Explorer and locate data from a saved file, such as a trace or log, and immediately load it into the Message Analyzer default **Analysis Grid** viewer.
  - **From Other File Sources** — click this item to display the **File Selector** dialog, from where you can specify input file sources that have a unique format. Currently, the **File Selector** is limited to working with Azure storage binary large objects (BLOBs) only. For more information, see Handling Azure Data.
- **Recent Files** — click this item to see a list of up to 10 recent files from which you loaded data into Message Analyzer. Data from a file in the **Recent Files** list is immediately loaded into Message Analyzer, as no further configuration is required.
- **Favorite Scenarios** — click this item to quickly start a Live Trace session with a single click on a **Trace Scenario** item in the list, for example, the **Local Network Interfaces**, **Loopback and Unencrypted IPSEC**, or **Pre-Encryption for HTTPS** scenario.
- **Save** — saves changes to a Data Retrieval Session such as bookmarks, comments, and time shifts. If you click this item after capturing data from a Live Trace Session that is not yet saved, the **Save/Export Session** dialog displays with several options for saving data.
- **Save As** — opens the **Save/Export Session** dialog that provides several save configuration options for the current set of trace results, which includes saving filtered, selected, or all messages in the data set.
- **Start Page** — click this item in the **File** menu to display the Message Analyzer **Start Page**.
- **Exit** — click this item to close Message Analyzer. If you have any unsaved changes, you will be prompted with the option to save them.

Session Menu

The global Message Analyzer **Session** menu provides access to the features described in the list that follows.

- **New Viewer** — click this item to open the **New Viewer** drop-down list, from where you can choose one of the built-in data viewers. You can also select a built-in **Layout** or a custom **Layout** that you created for the **Chart**, **Analysis Grid**, or **Grouping** viewers, for the current in-focus session. Data viewers that you can select from the list consist of the following:
  - **Analysis Grid**
  - **Grouping**
  - **Pattern Match**
  - **Gantt**
  - **Chart**
  - **Interaction**
  - **Message Summary Tiles**
  - **Message Summary Lists**
  - **PerfMon**
  - **Charts (Deprecated)**
- **Edit Session** — click this item to launch the **Edit Session** dialog, which displays the initial configuration of either a Data Retrieval Session or Live Trace Session, depending on the type of existing session that is

currently in focus. Enables you to edit the session configuration and then re-run it with your applied changes.

- **Reparse** — click this command to reparse the messages in the current set of trace results. For example, you could reparse a displayed set of trace results after adding an SSL certificate for decryption in the **Options** dialog, so that you can display the decrypted data.
- **Shift Time** — click this item to display a submenu with the following items:
  - **Shift Time** — click this item to display the **Shift Time** dialog, from where you can specify a time shift for message **Timestamps** in a set of session results, to accommodate for time zone changes or skewed clock values in a message collection that is comprised of multiple disparate sources.
  - **Remove All Time Shifts** — click this command to remove any previously specified time shifts from the current in-focus session.
- **Data Source Filter** — click this item to display a submenu with the following items for manipulating the display of data sources:
  - **Edit** — click this item to display the **Data Source Filter** dialog, from where you can select one or more data sources that exist in a set of trace results for which you want to view data. Data for any unselected data source will not display after you click **Apply** to exit the dialog.
  - **Apply** — this command is enabled only after you have created a data source configuration, applied it, and then removed it with the **Remove** command. Enables the **Remove** command to toggle into the enabled state. Working together with the **Remove** command, it enables you to alternately **Apply** and **Remove** the current data source configuration.
  - **Remove** — click this command to remove the current data source configuration that you specified in the **Data Source Filter** dialog. Enables the **Apply** command to toggle into the enabled state. Working together with the **Apply** command, it enables you to alternately **Remove** and **Apply** the current data source configuration.
- ***ViewerName*** — a placeholder for which the name changes depending on the type of session viewer that is currently in focus. For example, if the **Grouping** viewer is in focus, then the name of this **Session** menu item changes to **Grouping**. If an **Analysis Grid** session tab is currently in focus, the item name changes to **Analysis Grid**, and so on. Click the viewer name item in the **Session** menu to display a submenu with a set of commands that apply to the indicated data viewer only. Typically reproduces the commands that exist on the respective data viewer toolbar.

Tools Menu

The global Message Analyzer **Tools** menu provides access to the features described in the list that follows.

- **Windows** — utilize interactive **Tool Windows** that respond to message selection or session selection to provide additional message details. The **Tool Windows** that are available consist of the following:
  - Session Explorer Tool Window — monitor operational status and session statistics, and observe real-time progress indicators when loading, capturing, filtering, sorting, finding, grouping data, and applying sequence matching; navigate among different data viewers in various sessions; and select new data viewers from a context menu.
  - Message Details Tool Window — view field names and values for any message that you select in the **Analysis Grid**.
  - Message Data Tool Window — highlight hexadecimal values for any field that you select in the **Details** window or **Analysis Grid**, including payloads.
  - Field Data Tool Window — display the value of any field that you select in the **Details** window.

- Bookmarks Tool Window — mark one or more messages of interest, which includes adding links, attachments, and different colored flags.
- Comments Tool Window — quickly add basic comments to one or more messages.
- Diagnostics Tool Window — currently a preview feature that summarizes diagnosis errors and through selection, enables you to easily jump to a corresponding diagnosis message in the **Analysis Grid** viewer. You can also filter **Diagnostics Tool Window** columns to isolate specific column data.
- Message Stack Tool Window — display the message stack for any selected message row in the **Analysis Grid** viewer.
- Decryption Tool Window — display statistics, summary, and analysis information for a decryption session.
- Selection Tool Window — undo erroneous message selections or maintain the context of multiple message selection in the **Analysis Grid** viewer in a separate space that is independent of the grid selection, to facilitate ease of analysis.
- Compare Fields Tool Window — enables you to display the values of an identical selected field in two separate messages for comparison and analysis.
- Field Chooser Tool Window — specify additional message fields in the Analysis Grid Viewer when it is in focus, for deeper analysis of message data. Specify additional message field Groups in the Grouping Viewer, when it is in focus. Expands the scope of data presentation and further enhances data examination and troubleshooting. Also use **Field Chooser** for configuration tasks when creating **Pattern Expressions**, **Chart** formulas, **Unions**, and so on.
- Output Tool Window — display this tool window to monitor the Message Analyzer log file output for errors when loading modules.
- Window Layout — this list item provides access to built-in **Window Layouts** that organize the **Analysis Grid** viewer with different **Tool Windows** as preset configurations that enable you to customize your working analysis environment. The presets consist of the following:
  - **Simple**
  - **Simple with Field Data**
  - **Network**
  - **Multiple Sessions**
  - **Protocol Development**
  - **Advanced**

- **Add-Ins** — click this item to display a submenu with the **Compare (Preview)** tool. A preview feature that displays the **Session Comparison Utility**, which performs type and field check differences between two specified sessions.
- **Aliases** — click this item to display the **Aliases** drop-down list, from where you can select an **Alias** to apply to the currently in-focus session viewer tab. An **Alias** substitutes for cryptic or otherwise unfriendly field values that display in the Message Analyzer **Analysis Grid** viewer, for example an IPv6 address or a TCP port, for ease of analysis. Also edit and manage **Aliases** from this drop-down list.
- **Unions** — configure and manage **Unions** of two or more fields that have identical values but different names in different data sources. Enables the correlation of field values from such sources with a single field name in Message Analyzer, for ease of analysis.

- **Asset Manager** — click this item to display the **Asset Manager** dialog, from where you can manage Message Analyzer asset collections such as **Message Analyzer Chart View Layouts**, **Message Analyzer Color Rules**, **Message Analyzer Correlations**, **Message Analyzer Filters**, **Message Analyzer Profiles**, **OPN Parsers**, and so on. Enables you to download asset collections and auto-sync them for automatic updates provided by a Microsoft web service. Also enables you to create custom feeds where you can share Message Analyzer assets that you developed on your own.
- **Options** — click this item to open the **Options** dialog, which displays the following configuration tabs:
  - **General** tab — enables you to specify various default global settings for Message Analyzer, as follows:
    - **Live Trace Message Buffer** — provides settings that determine the rate at which packets are dropped when exceeding the ETW buffer limit.
    - **Text Log Files** — provides a drop-down menu that enables you to select a predefined default or custom configuration file for parsing text logs.
    - **Updates** — enables you to check for updates to your Message Analyzer installation, either manually or automatically.
  - **Display** tab — provides the controls for setting the format for date-time and Binary Values:
    - **Time Display** — enables you to specify the date-time format that Message Analyzer will use across all features that display time data.
    - **Binary Values** — enables you to specify different formats for the display of numeric data, which includes **ASCII**, **Hex**, and **Decimal** options.
  - **Decryption** tab — provides the controls that allow you to import and select server certificates and to specify passwords that are required to enable Message Analyzer to decrypt traffic that is encrypted with the Transport Layer Security (TLS) and Secure Sockets Layer (SSL) security protocols.
  - **Features** tab — provides for selection of preview features such as viewers and **Tool Windows** that you want to enable in Message Analyzer.
  - **Memory** tab — specifies the current memory statistics for Message Analyzer, the current state of **Server Garbage Collection**, and instructions for how to disable **Server Garbage Collection** to reduce memory consumption.
  - **Parsing** tab — enables you to reparse a set of trace results based on alternate ports that you specify for specific protocols, to accommodate for network traffic that used alternate ports for security purposes.
  - **Privacy** tab — enable Microsoft to collect information about stability issues, system configuration, and your frequently used features to help improve your Message Analyzer experience. Also enables you to opt-in to provide feedback on features, if you opted-out at the first Message Analyzer start up.
  - **Profiles** tab — enables you to select built-in **Profiles** that provide preset **Analysis Grid**, **Grouping**, and **Chart** viewer **Layouts** that create an interactive and integrated analysis environment for specific types of data associated with different file types. A selected **Profile** triggers when you load data from a specific file type, such as a .cap, .pcap, .etl, or .log file, with which the selected **Profile** is functionally associated. You can also create and edit your own custom **Profiles**.
  - **WPP** tab — provides a convenient interface that simplifies the input configuration for specifying symbol files that are required to parse WPP-generated events.

Help Menu

The global Message Analyzer **Help** menu provides access to the features described in the list that follows.

- **Feedback Center** — click this item to open the **Feedback Center** dialog, from where you can provide feedback for predefined questions about various Message Analyzer features. Note that the feedback features are reproduced by the **Feedback Center** and **Feedback** controls in the upper-right section of the Message Analyzer user interface.
- **Message Analyzer Feedback** — click this item to display a submenu that offers the following options for providing feedback:
  - **Send a Smile** — tell us what you liked.
  - **Send a Frown** — tell us what we can do better.
  - **Report a Bug** — provide us with details about problems that you encountered.
  - **Request a Feature**— request a feature that you think would improve your experiences with Message Analyzer.
- **View Help** — click this item to access the Message Analyzer Operating Guide on TechNet.
- **Message Analyzer Team Blog** — click this item to open the Message Analyzer Team Blog site to review regularly posted blog articles on various Message Analyzer subjects. Also enables you to provide comments, ask questions, and receive feedback directly from Microsoft.
- **Discussion Forum** — click this item to open the Message Analyzer Forum site, where you can start a discussion thread and view responses from the Message Analyzer community of users and Microsoft.
- **Check for Updates** — enables you to manually initiate a check for updates to your Message Analyzer installation.
- **About** — displays release information such as the current Message Analyzer version and build number, along with a Privacy Alert.

Global Toolbar

Message Analyzer also provides a global toolbar, from which you can access most of the same features that are provided in the global menus, but with fewer clicks and submenu selections. Fuller descriptions for many of these features are provided in the Global Menus section. The global Message Analyzer toolbar contains the following items:

- **New Session** — click this toolbar item to open the **New Session** dialog, where you can choose a source from which to acquire data, as described earlier.
- **Favorite Scenarios** — click this toolbar item to quickly start a Live Trace session with a single click on a **Trace Scenario** item in the list.
- **Open** — provides a submenu that enables you to display either the **Open** dialog for Windows Explorer, or the **File Selector**, which is currently configured to open Azure storage BLOBs only.
- **Save** — click this toolbar button to save bookmarks, comments, and time shifts, as described earlier.
- **New Viewer** — click this toolbar drop-down list to open additional data viewers against a set of live trace results or loaded data, for diagnostic and analysis purposes. For example, you might select the **Grouping** viewer or **Chart** viewer with a chosen **Layout** for your selection.
- **Edit Session** — click this toolbar button to open the **Edit Session** dialog to reconfigure an existing Data Retrieval Session or Live Trace Session, and then apply your changes by clicking the **Apply** button in the dialog. Note that the session will reload all data if the changes you make are in **Full Edit** mode, otherwise, you can add data from more files without incurring a reload of existing data.

- **Restart** — this toolbar button enables you to restart a stopped Live Trace Session, which starts a new Live Trace Session and causes you to lose any previous data that you collected. You can also restart a stopped Live Trace Session that you edited. For example, after you edit an existing Live Trace Session through the **Edit Session** dialog, click the **Restart** button to apply the configuration changes that you specified. Note that this action also starts a new Live Trace Session and abandons any previous data you collected. Also note that the **Restart** control is included on the global toolbar only if you are displaying a set of trace results from a Live Trace Session.
- **Pause/Resume** — click this toolbar button to pause a Live Trace Session in progress and then click it again to resume capturing data again. Toggles back and forth between the paused state and the capture resumed state as you click this button successively. Note that this control is included on the global toolbar only if you are displaying a set of trace results from a Live Trace Session.
- **Stop** — click this toolbar button to stop the capture of messages in a Live Trace Session. Note that this control is included on the global toolbar only if you are displaying a set of trace results from a Live Trace Session.
- **Shift Time** — specify time shifts that enable you to adjust the time stamps in a message set, for example to compensate for machine skew or time-zone changes across multiple data sources.
- **Aliases** — click this toolbar item to display the **Aliases** drop-down list, from where you can select a one of the default **Aliases** or a custom **Alias** that you created, and apply it to data in the currently in-focus session viewer.
- **New Union** — click this toolbar item to display the **Edit Union** dialog, from where you can create a **Union** of two or more fields that have identical values but different names in different data sources.
- **Window Layout** — click this toolbar item to display a list of **Window Layouts** from which you can select a preset **Tool Window** layout to accompany the **Analysis Grid** viewer. Enables you to create a custom working environment that accommodates the type of analysis that you typically perform.

Start Page

The Message Analyzer **Start Page** contains a top row of button commands for quick access to the following features or functions:

- **New Session** — to start a new Data Retrieval Session or a Live Trace Session, open the **New Session** dialog with a single click on the **New Session** button.
- **Start Local Trace** — start a local Link Layer trace in a basic configuration with a single click, with no additional configuration required.
- **Open** — launch the Windows Explorer **Open** dialog with a single click, to locate a saved file containing message data and quickly load it into Message Analyzer.
- **Discussion/Voting Forum** — go to the Message Analyzer Forum to start a discussion thread, where you can get feedback from Microsoft. Also vote on discussions that you find helpful.
- **Message Analyzer Team Blog** — go to the Message Analyzer Blog to review numerous Blog postings about Message Analyzer features and use. Also, leave comments, rate articles, and get feedback from Microsoft. In addition, you will find a link to Message Analyzer training Videos on this site under **Important Links**.

From the **Start Page**, you can also view **Recent Files**, **Favorite Scenarios**, edit **Favorite Scenarios**, and review **News** items.

Layouts for the Chart Viewer

Message Analyzer enables you to select various built-in **Chart Viewer Layouts** against a set of trace results to enhance your data analysis perspectives. Message Analyzer also provides the configuration tools and other features needed to create, edit, save, and share **Layouts** for the **Chart** viewer that you can configure with custom **Bar** element, **Pie** slice, **Timeline** graphic, and **Table** grid visualizer components. Note that each of these components are used in the built-in **Protocol Dashboard** viewer, which is accessible from the **Charts (Deprecated)** drop-down list in the **New Viewer** drop-down on the global Message Analyzer toolbar.

When creating new **Layouts** for the **Chart** viewer, you can use the **Field Chooser Tool Window** to specify message fields for your **Layout** so that you can monitor values and you can also configure data manipulation formulas to create diverse top-level data summary configurations and statistical displays that empower visual analysis capabilities. Note that Message Analyzer **Chart** viewer **Layout** configurations also support **Unions** and **Union** sets. You can access the configuration controls for creating new **Layouts** by clicking the **Edit** item in the **Chart** drop-down list that appears on the global Message Analyzer **Session** menu whenever a **Chart** is displayed and currently in focus.

Sharing Infrastructure

Message Analyzer provides a **Sharing Infrastructure** that enables you to download default user asset collections to your local Libraries for manipulating and viewing data; this also includes downloading OPN packages for parsing messages that you capture with Message Analyzer. You can synchronize these collections and packages for automatic updates that are periodically pushed out by a Microsoft web service to the default **Message Analyzer** subscriber feed. You can view this feed from the **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu. Because the user **Libraries** are integrated with the Sharing Infrastructure, you can import, export, and share these items with others, including any that you create or modify. Library item types include **Message Analyzer Trace Scenarios**, **Message Analyzer Filters**, **Message Analyzer Chart View Layouts**, **Message Analyzer Color Rules**, **Message Analyzer Profiles**, **Message Analyzer Viewpoints**, and so on. To enable sharing these Library items, you can configure your own user feeds or post items to a user file share. You can also manage all user Library types with the common and centralized management dialog.

Other Capabilities

Other prominent Message Analyzer capabilities include the following:

- **Configuring a Remote Capture** — capture remote traffic on Windows 8.1, Windows Server 2012 R2, and Windows 10 hosts. Use the **Remote Network Interfaces Trace Scenario** with the **Microsoft-Windows-NDIS-PacketCapture** provider to target one or more computers with supported operating systems for remote capture in a Live Trace Session.
- **Configuring Host Adapter and Hyper-V-Switch Filters** — capture traffic from one or more host adapters and/or virtual machines (VMs) that are serviced by a Hyper-V Switch on remote Windows 8.1, Windows Server 2012 R2, or Windows 10 hosts, or on the local computer. Customize the capture configuration by specifying packet traversal paths on switch extension layers and on the NDIS driver filter stack, along with configuring other special filters, such as packet **Truncation**, **EtherType**, and **IP Protocol Number** filters through use of the **Advanced Settings - Microsoft-Windows-NDIS PacketCapture** dialog.
- **Processing MOF-Generated Events** — fully parse messages that are captured by Message Analyzer from MOF-instrumented providers. Message Analyzer supports registered event providers on your system that use the MOF schema as the basis of generating their events.
- **Processing WPP-Generated Events** — parse and display Windows software trace preprocessor (WPP)-generated events. Because such events make use of the ETW framework, Message Analyzer can capture them live or load them from a saved event trace log (ETL) file. To enable parsing of WPP-generated events, users must provide supplementary information that defines the WPP event structure.

- **PEF-WFP Fast Filters** — specify **Fast Filters** for the **Microsoft-PEF-WFP-MessageProvider** in a **Loopback and Unencrypted IPSEC** trace.
- **PEF-NDIS Fast Filters** — configure logically chained **Fast Filter** groups that you assign to host adapters by using the **Advanced Settings - Microsoft-PEF-NDIS PacketCapture** dialog in a **Local Network Interfaces** trace on Windows 8 and earlier hosts.
- **Microsoft-Windows-NDIS-PacketCapture Provider** — enabled for capturing remote traffic, with support for capturing from VMs that are managed by a Hyper-V-Switch. Also supports advanced filtering that includes packet direction, NDIS stack, and Hyper-V extension layer filters.
- **Microsoft-PEF-WFP-MessageProvider** — since Message Analyzer v1.3, the **Microsoft-PEF-WFP-MessageProvider** has had the capability to capture messages from remote computers that are running the Windows 10 operating system. You can capture this data in any **Trace Scenario** that uses this provider by starting a Live Trace Session from any computer that is running the Windows 8.1, Windows Server 2012 R2, or the Windows 10 operating system.
- **Capturing in Promiscuous Mode (P-Mode)** — configure the **Microsoft-Windows-NDIS-PacketCapture** provider to capture local or remote traffic on network adapters that support p-mode by simply selecting supporting adapters in the interface selection list of the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture dialog.
- **Filtering Language** — discover how to write your own Filter Expressions for filtering data that is loaded into Message Analyzer, captured live, or analyzed after trace results are complete.
- **ResponseTime** — add this **Global Annotation** entity from the **Field Chooser** as a data column in the **Analysis Grid** viewer. Enables you to measure the time interval between a request operation to a server and the first server response, to provide a context for assessing server performance.
- **OPN Definitions** — display OPN definitions for capture modules or message fields from the **Analysis Grid** viewer or **Details** window context menu, respectively.
- **Analysis Grid Toolbar Features** — the **Analysis Grid** viewer now provides a toolbar for quick access to tools that assist in common analysis tasks.
- **Grouping Viewer** — select the **Grouping** viewer to organize your traffic into summary group hierarchies that expose targeted information that you can quickly extract from a large data set, which can otherwise be difficult to achieve.
- **Track Fields and Properties** — a **Details** window feature that exposes the values of message-specific fields, along with the additional global message properties and global annotations that are generated by Message Analyzer, to enable you to do the following from the **Details** window context menu:
  - Utilize these new entities when configuring filters, groupings, adding columns, and so on.
  - Track any message field, global property, or global annotation value — this enables you to compare the value of the same field or property in other messages that you select or scroll to in the **Analysis Grid** viewer.
- **Pattern Match Viewer** — enables you to execute predefined **Pattern** expressions that locate message sequences or patterns that occur across a set of trace results, for example, a **TCP Three-Way Handshake** pattern. Also provides facilities for creating your own **Pattern** expressions, with or without user interface automation assistance.
- **Parse As** — a global **Options** dialog feature that enables you to parse an existing set of trace results with a different port value, for example, one that deviates from a standard port value for a particular protocol. Accommodates traces that used alternate port values for security purposes.

Important for Network Monitor Users

Microsoft Message Analyzer dramatically extends the network traffic diagnostics and analysis capabilities of Microsoft Network Monitor. This is enabled by a diverse set of data viewers and graphical view **Layouts** that create high-level data summaries; top-level messages and Operations with message stack and fragment encapsulation; **Viewpoints**, **Profiles**, and **Diagnosis** messages; and so on. In addition, Message Analyzer implements former Network Monitor features that enable you to do the following:

- Correlate process name and process ID data natively in Message Analyzer.
- Simulate the Network Monitor message display with the Message Analyzer **Flat Message List** feature.
- Simulate the Network Monitor **Network Conversations** tree with the Message Analyzer **Grouping** viewer.

For a high-level comparison of Message Analyzer and Network Monitor features and why new approaches have been taken for capturing, displaying, and analyzing message traffic, see the following Message Analyzer Blog articles.

Message Analyzer v1.3 vs Network Monitor v3.4 Message Analyzer: Why so different from Network Monitor?

Quick Session Startup

5 minutes to read

If you have not yet started Message Analyzer, see Starting Message Analyzer for the First Time. After you start Message Analyzer, the first user interface that displays is the **Start Page**. This page is designed to enable you to acquire input data very quickly, for example, through a Data Retrieval Session or a Live Trace Session.

> **IMPORTANT**
>
> To start a Live Trace Session immediately with a single action and no further configuration, click the **Start Local Trace** button on the **Start Page**, as shown in the figure below, to begin capturing live traffic at the Link Layer with the **Local Network Interfaces Trace Scenario**. Note that you should be running Message Analyzer as an Administrator for this scenario to function properly.

Other methods for quick access to input data are described in Performing a Live Capture.

Start Page Organization For Quick Access to Input Data

The **Start Page** is organized into an upper row of action buttons that provide fast access to input message data, and below that are recent files and scenario lists that likewise also provide quick access to input data. Input message data consists of message traffic that you capture live or saved data that you import into Message Analyzer. The **Start Page** facilitates the acquisition of this input data by enabling you to get started very quickly with a Live Trace Session or a Data Retrieval Session, where you capture network traffic live or load data from saved files or logs into Message Analyzer, respectively. The Message Analyzer **Start Page** is shown in the following figure:

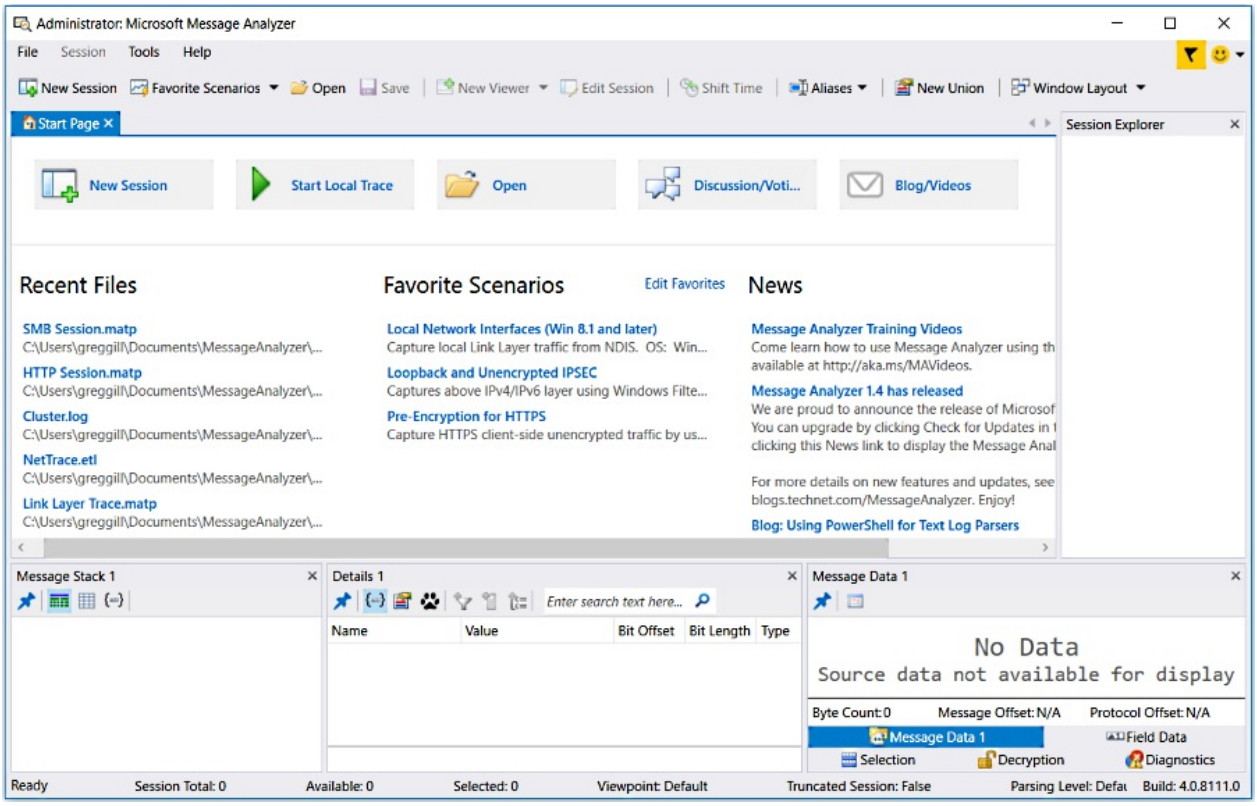

Figure 1: Message Analyzer Start Page

The following provides details about the **Start Page** features that enable quick access to input data:

- **New Session** button — contained in the upper row of action buttons. Click this control to start the configuration phase of a Live Trace Session or Data Retrieval Session, by displaying the **New Session** dialog. This dialog contains configuration screens that enable you to quickly access the settings and components that you want to specify for a live trace or data import, for example, a predefined **Trace Scenario**, ETW Provider, Filters, a session viewer, target host/s, saved traces or logs, and so on.

  For additional information about Message Analyzer sessions, see Starting a Message Analyzer Session.

- **Start Local Trace** button — contained in the upper row of action buttons. Click this control to immediately start a local Link Layer trace while using the **NDIS-PacketCapture** provider with no additional configuration required, although you need to run Message Analyzer with the **Run as administrator** startup option in this case. This trace captures network traffic with either the **Microsoft-PEF-NDIS-PacketCapture** provider or the **Microsoft-Windows-NDIS-PacketCapture** provider, depending on which operating system your computer is running, as described in Built-In Trace Scenarios. When you use this feature to capture live data, Message Analyzer automatically creates a Live Trace Session and displays the results in the **Analysis Grid** viewer or in the viewer that is set as the default in the **Default Profile** pane on the **Profiles** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu.

  For additional details, see Performing a Live Capture.

- **Open** button — contained in the upper row of action buttons. Click this control to load data into Message Analyzer from saved trace or log files that are supported by Message Analyzer, as described in Locating Supported Input Data File Types. When you click the **Open** button on the **Start Page**, the Windows **Open** dialog displays, from where you can navigate to a saved trace or log file. By default, this dialog opens to the `Documents\MessageAnalyzer\Traces\` folder on your computer; however, if you change the location for storing your files, the **Open** dialog will persist that location.

  After you select one or more files containing the data you want to load, Message Analyzer automatically creates a Data Retrieval Session configuration in the **New Session** dialog, which lists the files you selected. At this point, prior to starting the Data Retrieval Session, you have the option to configure a **Time Filter**, **Session Filter**, and/or a **Parsing Level**, and you can also select a data viewer of choice. Note that when you load data into Message Analyzer through the **Open** feature, the resulting message set will be chronologically ordered in the **Analysis Grid** viewer (assuming this is the data viewer you are using).

  For related information, see Performing Data Retrieval.

- **Recent Files** list — contains a list of trace and log files that you recently saved or accessed from Message Analyzer. This list displays a maximum of 10 files and enables you to very quickly resume work from a previous Analysis Session. When you click a trace or log file in the list, Message Analyzer immediately loads the data and displays it in the default data viewer, such as the **Analysis Grid**. In addition, Message Analyzer automatically creates a Data Retrieval Session that you can edit later on by clicking the **Edit Session** button on the global Message Analyzer toolbar. You might do this to alter the data display by specifying a **Time Filter**, **Session Filter**, or **Parsing Level**.

- **Favorite Scenarios** list — contains a list of **Trace Scenarios** that are set to Favorite status. By default, the list contains the predefined **Local Network Interfaces**, **Loopback and Unencrypted IPSEC**, and **Pre-Encryption for HTTPS** scenarios. When you click one of the **Trace Scenarios** in this list, Message Analyzer automatically starts a local Live Trace Session with no additional configuration required, at which point Message Analyzer begins to capture data in accordance with the scenario's built-in design. After stopping the session, you have the option to edit it as previously described.

  Note that you can add other predefined **Trace Scenarios** or custom scenarios of your own to the **Favorite Scenarios** list, by clicking the **Edit Favorites** link above the list. This action displays the **Edit Favorites** dialog that enables you to set a new Favorite or undo an existing one. To configure a Favorite, click the white-colored star to the left of the scenario in the **Edit Favorites** dialog, at which time the star

changes to the color yellow and the scenario is added to the **Favorite Scenarios** list on the **Start Page** and to the **Favorites** category of the **Message Analyzer Trace Scenarios** asset collection. To undo a Favorite, simply click the yellow star to remove the scenario from the **Favorite Scenarios** list and the **Favorites** category of the asset collection.

> **NOTE**
>
> The **Message Analyzer Trace Scenarios** asset collection also appears in the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog. You can also set Favorites from this location.

More Information

To learn more about the built-in **Trace Scenarios** provided with Message Analyzer, see the [Built-In Trace Scenarios] (built-in-trace-scenarios.md) topic.

Technology Tutorials

2 minutes to read

This section contains several tutorials to help you better understand Message Analyzer concepts, starting with a functional overview of Message Analyzer and concluding with conceptual overviews of the architectures and frameworks upon which it is based. The complexities of the underlying systems are briefly described in these tutorials so that readers will have a greater comprehension of Message Analyzer capabilities. The technology tutorials are contained in the topics below.

Message Analyzer Tutorial
PEF Architecture Tutorial
ETW Framework Conceptual Tutorial

Message Analyzer Tutorial

103 minutes to read

This section begins with some background concepts about Microsoft Message Analyzer and then goes into several mini-tutorials or Getting Started Primers that will help you get started with using this unique tool. Links are provided throughout so that you can navigate to more information about the described features as needed.

Go To Procedures

To go directly to procedures that provide examples of using Message Analyzer, see the following topics:

Procedures: Quick Start
Procedures: Using the Network Tracing Features
Procedures: Using the Data Retrieval Features
Procedures: Using the Data Viewing Features
Procedures: Using the Data Filtering Features
Procedures: Using the Asset Management Features
Procedures: Using the Chart Viewer Layout Configuration Features

---

Introduction

The overarching and new approach that Message Analyzer uses when capturing traffic is to limit network noise and to expose at top-level both the issues that occur at lower levels and hidden information that is critical to quick analysis. Message Analyzer does this by the following.

- Enabling you to remove lower-layer messages in a capture so you can focus on higher-layer data of interest.
- Displaying individual message summaries as well as high-level overviews of trace statistics and trends.
- Exposing diagnostics data in top-level transactions.
- Creating top-level Operation nodes that encapsulate request and response messages for quick assessment of details, such as server response time.
- Locating message fragment reassemblies within the origins tree (stack messages) rather than in a dispersed chronological display.
- Enabling you to control the layer up to which Message Analyzer will parse, with the use of Parsing Levels.
- Enabling you to "select" specific data that you want to view through filtering.

In this manner, the important information that you need to see for any particular message is readily exposed at top-level in the **Analysis Grid** viewer, which is the main analysis surface that Message Analyzer provides.

Another significant feature that enables you to focus on messages of interest is **Viewpoints**, which display data from the perspective of a chosen protocol, module, or layer with no messages above it. For example, you could select a TCP **Viewpoint** and drive all TCP messages to top-level in the **Analysis Grid** to facilitate better analysis of TCP messages. This is in contrast to Message Analyzer's predecessor Network Monitor, which shows only flat or static message packets in original capture order and does not hide any noise, reassemble fragments, or simulate protocol behavior to allow for interpreting states and maintaining a protocol model, such as Message Analyzer does. Moreover, Message Analyzer formalizes its parser definitions to enable more artifacts to be derived from them, such as test cases and documentation.

You will learn more about these features in the next few sections that provide an overview of acquiring data

through a Message Analyzer session and using various tools to focus data capture and analysis on specific types of data. After these sections, you can review the Getting Started Primers.

Acquiring Data Through a Message Analyzer Session

Message Analyzer enables you to capture, display, and analyze protocol messaging traffic, and to trace and assess system events, Windows component events, and device messages. It also provides the capability to retrieve, aggregate, and analyze data from one or more saved traces, which includes support for the .etl, .cap, .pcap, .pcapng, .tsv/.csv, .evtx, and .log input file formats, in addition to Message Analyzer native files in the .matp or .matu format, as described in Locating Supported Input Data File Types. If you work with text based .log files, Message Analyzer enables you to retrieve data from various common text .log file types with the use of built-in text log parsers that are described in Parsing Input Text Log Files. Also note that if you have a custom text .log file, an extensibility feature of the Microsoft Protocol Engineering Framework (PEF) enables Message Analyzer to retrieve its data with the use of a custom configuration file. However, you will need to create this file in order to fully parse your text log, as described in Parsing Input Text Log Files. Message Analyzer also enables you to extend the functionality of the **Chart** viewer by creating custom view **Layouts** of your own design, as described in Extending Message Analyzer Data Viewing Capabilities.

Message Analyzer makes use of two different types of sessions to acquire input data, as described in Starting a Message Analyzer Session. These consist of a Live Trace Session and a Data Retrieval Session, which provide data from the live capture of network traffic, events, system messages, and device messages; and saved traces, logs, and text logs, respectively. In a Live Trace Session, PEF provider-drivers and/or other system ETW Providers listen for and capture protocol messages and events at various stack layers or from other components. The messages and events are passed to the PEF Runtime where they are decoded by Open Protocol Notation (OPN) parsers and then temporarily saved in a Message Store. To access and display these messages, Message Analyzer consumes the PEF Runtime data, as described in the PEF Architecture Tutorial. Messages are displayed by default in the **Analysis Grid** viewer, where you can begin your data analysis process; however, other data viewers and various **Tool Windows** are also available to streamline message analysis.

Live Trace Session

In a Live Trace Session, you have the option to capture data from the local computer and/or multiple remote computers in concurrent subsessions that return all data to the common initiating live session that you configure with a chosen data viewer. Moreover, the local computer is the default host on which a Live Trace Session captures data; however, if you specify valid connection/authentication credentials for other remote computers, you can capture data simultaneously on those computers as well. Message Analyzer also provides you with the flexibility to run multiple concurrent Live Trace Sessions, optionally with each having different message provider and filtering configurations, to target different computers. You can do this by simply adding one or more **Live Trace** data sources in the **New Session** dialog, specifying the hosts from which to capture the data, and selecting or creating **Session Filters**, as described in Configuring Session Scenarios with Selected Data Sources.

> **TIP**
>
> **Quick Tracing** — to get started very quickly with a Live Trace Session, you can make use of **Start Page** features that enable you to start a new Local trace session at Link Layer or begin the configuration phase for a new session—with a single click—as described in Quick Session Startup.

More Information

To learn more about configuring a Live Trace Session, see Capturing Message Data.

Data Retrieval Session

In a Data Retrieval Session, Message Analyzer enables you to retrieve and aggregate saved message collections from multiple sources, including traces and logs, in any combination. This means you can mix and merge data from any of these sources and display it in the **Analysis Grid** or other selected data viewer. If you know that

certain events of interest have occurred at a particular time in a collection of data sources, you can configure a **Time Filter** to view data in a window of time that you specify to eliminate extraneous data and improve performance. You can also set **Time Shifts** to accommodate for different time zones or skewed machine times across different data sources. You might also select a built-in **Session Filter** or configure one of your own design to return specific data that is based on the filtering criteria that you specify, while at the same time further improving performance.

Special Input Sources
Message Analyzer also provides access to special input sources such as Azure Storage Tables, Azure Storage Blobs, Event Logs, SQL databases, and Operations Management Suite (OMS) logs. It also provides an interface from where you can write PowerShell queries. For access to most of these input sources, you will need authentication credentials. The user interface for all of these input sources is located in the **New Session** dialog, which is accessible from the **Start Page** by clicking the **New Session** button.

Message Analyzer also provides a set of built-in parsers for common text logs such as Cluster, Netlogon, IIS, and so on. In addition, if you have a proprietary text log with a unique format, you have the option to create an OPN configuration file which enables Message Analyzer to parse the data in your log file, as described in Parsing Input Text Log Files.

More Information
To learn more about configuring a Data Retrieval Session, see Retrieving Message Data.
To learn more about accessing data from the previously mentioned special input sources, see Acquiring Data From Other Input Sources.

Focused Tracing and Analysis

Although Message Analyzer enables you to capture messages from many system components, the PEF providers used by Message Analyzer enable you to capture data at several different layers, which provide unique inspection points into the protocol stack. For example, by specifying any **Trace Scenario** that uses the **Microsoft-PEF-WFP-MessageProvider**, you can focus on capturing messages above the IP/Network Layer by filtering out lower-level Link Layer messages through the Windows Filtering Platform (WFP), upon which the **Microsoft-PEF-WFP-MessageProvider** is based. Moreover, by specifying any **Trace Scenario** that uses the **Microsoft-PEF-NDIS-PacketCapture** or **Microsoft-Windows-NDIS-PacketCapture** provider, you can capture messages at Link Layer and above. Message Analyzer also enables you to temporarily set a predefined **Viewpoint** that filters, reorganizes, and redisplays the data from the perspective of a selected protocol or module type, such as HTTP, TCP, SMB, or ETW, so that you can focus on specific message traffic that is defined by the **Viewpoint**, while removing all messages above the **Viewpoint** level to create a focused set of messages.

You can also select a predefined **Parsing Level** that controls the stack level to which Message Analyzer parses, while passing certain messages in these scenarios that are useful to your data analysis perspective, as described in Setting the Session Parsing Level. In addition, you can make use of **Aliases**, as described in Using and Managing Message Analyzer Aliases, to configure user-friendly names for cryptic field values; and you can take advantage of the **Unions** feature, described in Configuring and Managing Message Analyzer Unions, to correlate differently named fields that are of the same type in different data sources. You can even capture and analyze loopback traffic for local application communications that use the IPv4 or IPv6 loopback addresses, by specifying the **Loopback and Unencrypted IPSEC** or **Local Loopback Network Trace Scenario**, as described in Built-In Trace Scenarios.

You also have the option to *select* specific data that you want to isolate for focused analysis by making use of any of the following:

- **Fast Filter** — a provider/driver-level filter that is very fast and efficient, as described in PEF-NDIS Fast Filters.

- **Keyword Filter** — returns only the events from an ETW Provider that are defined by one or more designated event **Keywords**, as described in System ETW Provider Event Keyword/Level Settings.
- **Session Filter** — creates a focused set of trace results that is determined by filtering criteria, as described in Working with Session Filters in a Live Trace Session.
- **Time Filter** — creates a window of time in which to view data, as described in Applying an Input Time Filter to a Data Retrieval Session and Applying a Time Filter to Session Results.
- **View Filter** — when applied to a set of trace results, passes only the message data that meets the filtering criteria that you specify, as described in Applying and Managing Filters. Enables you to create a focused set of results during an Analysis Session.

Furthermore, Message Analyzer enables you to decrypt data that is encrypted with the Transport Layer Security (TLS) and Secure Sockets Layer (SSL) protocols, for example Remote Desktop Protocol (RDP) and HTTPS messages, respectively. The **Decryption** feature also provides a **Decryption Tool Window** that presents summary and statistical data for the decryption session to facilitate analysis, as described in Decrypting TLS and SSL Encrypted Data.

These capabilities solve many inherent capture, data display, and analysis problems, such as the visibility of encrypted data, assessment of loopback traffic that is enabled by the **Local Loopback Network** scenario, and seeing traffic from the **Viewpoint** of a protocol. The underlying technologies that support Message Analyzer also machine-validate message structure and values, behavior, and architecture based on protocol specifications; and if errors occur, they are surfaced very quickly to top-level as Diagnosis messages. To this end, Message Analyzer also provides a **Diagnostics Tool Window** that summarizes all the Diagnostic messages in a trace, which interactively drives selection of corresponding messages in the **Analysis Grid** viewer to facilitate further review of message **Details**, **Message Stack** information, and **Message Data**.

> **NOTE**
> Message Analyzer is also an effective tool for testing and verifying protocol implementations. See the Open Specifications documentation library for more information about protocol technical specifications.

Getting Started Primers

The sections that follow provide brief conceptual tutorials that serve as getting started primers for Message Analyzer functionality. These tutorials correspond to the major tasks that you perform from the Message Analyzer user interface, where you can:

Capture Message Data
Retrieve Message Data
Edit Message Data
View Message Data
Filter Message Data
Analyze Message Data
Save Message Data

Capture Message Data

When capturing data live, Message Analyzer makes use of various message providers that focus on different layers or types of data. These providers are included in every Message Analyzer installation and consist of common Microsoft-PEF providers, the Microsoft-Windows-NDIS-PacketCapture provider, and various ETW Providers that are registered on the Windows system by default. These providers are briefly described in "Common Message Providers Used by Message Analyzer", which follows. Thereafter, this section describes how

to configure and start a Live Trace Session; provides examples of the Message Analyzer global options you can set; describes how Message Analyzer integrates event tracing into the capture process; how to optimize ETW sessions; and how Message Analyzer parses messages from MOF-based system ETW providers. The subject matter is discussed in the following topics.

Configuring a Live Trace Session
Starting a Live Trace Session
Setting Message Analyzer Global Options
Protocol Modules and Specifications
Integrating Event Tracing
Optimizing ETW Session Performance
Using MOF-Based ETW Providers

Common Message Providers Used by Message Analyzer The following message providers are included in Message Analyzer **Trace Scenarios**, which contain either one of these providers as the exclusive data source or a combination of several providers, depending on the scenario requirements.

- **Common Microsoft PEF Message Provider-Drivers** — all PEF drivers are instrumented with Event Tracing for Windows (ETW) provider technology, which enables them to take advantage of the ETW event tracing, buffering, logging, and event delivery infrastructure. In addition to numerous system ETW providers and other message capture components, all Message Analyzer installations contain the PEF provider-drivers in the list that follows, the configurations for which are accessible after you select a **Trace Scenario** from the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog for a Live Trace Session.

  > **IMPORTANT**
  > Some of the message providers described in this section may be different than what you find on your computer, because of an operating system version dependency. For example, on computers running the Windows 7, Windows 8, or Windows Server 2012 operating system, the **Microsoft-Windows-NDIS-PacketCapture** provider does not exist for the **Local Network Interfaces Trace Scenario**. Instead, the **Microsoft-PEF-NDIS-PacketCapture** provider is included in the **Local Network Interfaces** scenario on those computers. On computers running the Windows 8.1, Windows Server 2012 R2, and Windows 10 operating systems, the **Microsoft-Windows-NDIS-PacketCapture** provider is installed as part of the operating system and is used in the **Local Network Interfaces**, **Remote Network Interfaces**, and other **Trace Scenarios**.

  - **Microsoft-PEF-NDIS-PacketCapture** provider — an ETW-instrumented, Network Data Interface Specification (NDIS) light weight filter (LWF) driver that captures Ethernet frames at the Link Layer and delivers them to Message Analyzer through the ETW infrastructure. Also includes the capability to configure **Fast Filters** that operate efficiently at the driver-level to isolate specific message types, thereby passing less data and reducing system loads and resource consumption.

    More Information
    To learn more about the **Microsoft-PEF-NDIS-PacketCapture** provider, see [Microsoft-PEF-NDIS-PacketCapture Provider] microsoft-pef-ndis-packetcapture-provider.md).

  - **Microsoft-PEF-WFP-MessageProvider** — an ETW-instrumented driver that is based on the Windows Filtering Platform (WFP). It captures message traffic above the IP/Network Layer and delivers that traffic to Message Analyzer through the ETW infrastructure. This provider also enables you to configure **Fast Filters** to isolate specific messages of interest and improve trace performance. This provider is now enabled for remote capabilities when capturing data on remote Windows 10 computers only. In addition, you can set the **Select Discarded Packet Events** option when configuring this provider to log discarded packets.

More Information
To learn more about the **Microsoft-PEF-WFP-MessageProvider**, see [Microsoft-PEF-WFP-MessageProvider] microsoft-pef-wfp-messageprovider.md).

- **Microsoft-PEF-WebProxy** — an ETW-instrumented provider that uses the Fiddler API and acts as an HTTP proxy to intercept and capture all HTTP traffic to and from a client web browser in unencrypted format. Also provides the capability to configure driver-level **Hostname** and **Port** filters to isolate specific messages and improve performance.

More Information
To learn more about the **Microsoft-PEF-WebProxy** provider, see Microsoft-PEF-WebProxy Provider.

- **Microsoft-Windows-NDIS-PacketCapture** provider — an ETW-instrumented provider that has remote capabilities along with special NDIS stack and Hyper-V-Switch extension layer filtering, adapter configurations, packet traversal path directivity, and other filters and specifiers that you can configure.

> **NOTE**
> The **Microsoft-Windows-NDIS-PacketCapture** provider with remote capabilities is used on the Windows 8.1, Windows Server 2012 R2, and Windows 10 operating system only, as described in Built-In Trace Scenarios.

More Information
To learn more about the **Microsoft-Windows-NDIS-PacketCapture** provider, see [Microsoft-Windows-NDIS-PacketCapture Provider] (microsoft-windows-ndis-packetcapture-provider.md).
To learn more about capturing messages from one or more remote hosts and configuring the **Microsoft-Windows-NDIS-PacketCapture** provider, see Configuring a Remote Capture.

- **System ETW providers** — write events for various components on your system that have been instrumented as ETW event providers. This includes ETW providers that define their events with the use of the following:
  - Standard provider manifests, as described in Event Manifest.
  - Managed object format (MOF) schemas, as described in Using MOF-Based ETW Providers.
  - Windows WPP-generated events that are issued by software components, as described in Loading WPP-Generated Events.

Configuring a Live Trace Session

You can specify the message providers that you want to use to capture data from the network or other components by configuring a Live Trace Session, as shown in the figure that follows. In the figure, the **Microsoft-PEF-WFP-MessageProvider** appears in the list after selecting the **Loopback and Unencrypted IPSEC Trace Scenario** in the **Select Scenario** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog. The **Windows-Firewall-Service** ETW Provider appears in the list after selecting this provider in the **Add System Providers** dialog that displays when you click the **Add Providers** drop-down list on the same toolbar and then select the **Add System Providers** item.

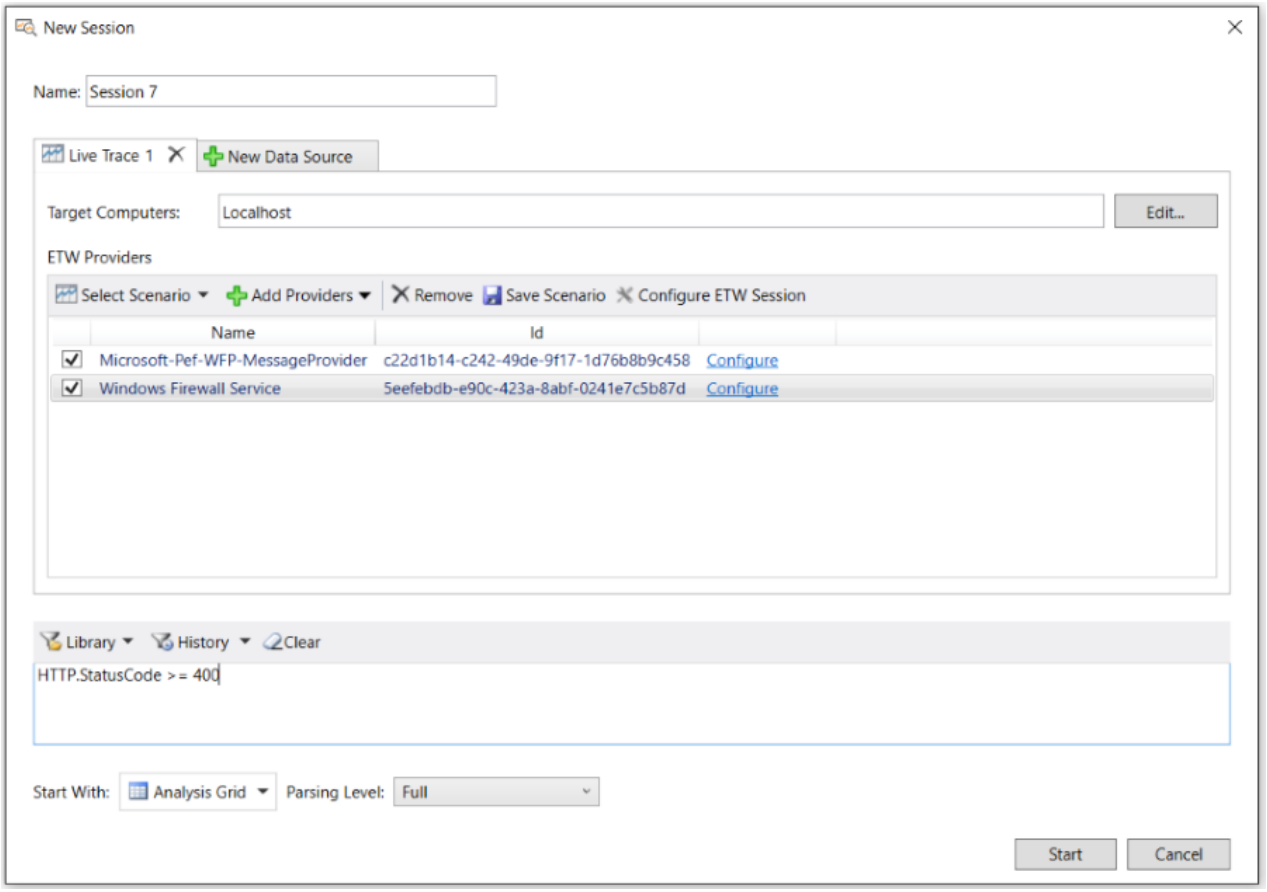

Figure 2: Message Analyzer Live Trace Session configuration

Predefined provider configurations are contained in all the built-in **Trace Scenarios** that you can select from the **Select Scenario** drop-down list on the **ETW Providers** toolbar on the **Live Trace** tab of the **New Session** dialog. These **Trace Scenarios** are templates that contain predefined message provider configurations that are tailored for capturing data from various components and/or at different stack layers.

Optionally, you can enhance the scope of data capture by adding other system ETW providers to the scenario. Also, if you have created and saved any custom **Trace Scenarios** by using the **Save Scenario** feature on the **ETW Providers** toolbar, these are also available for selection in the **My Items** category of the **Select Scenario** drop-down list. See Creating and Managing Custom Trace Scenarios for further details on creating your own scenarios. You can also modify the capture configuration of PEF and other **ETW Providers** from the **Live Trace** tab of a **New Session** to isolate specific message traffic and realize performance enhancements.

For example, by clicking the **Configure** link for a selected message provider in the **ETW Providers** list, such as the **Microsoft-PEF-WFP-MessageProvider**, you can display a configuration dialog and specify **Fast Filters** that work very efficiently at the kernel level. These low-level filters enable you to quickly retrieve specific messages that meet the filtering criteria that you specify, which reduces the scope of the data to be returned by the trace. In turn, this accelerates the data capture process and minimizes the Message Analyzer parsing time.

You also have the option to select or create a **Session Filter** for a Live Trace Session (or a Data Retrieval Session) to reduce the scope and count of messages that you retrieve, and as a result realize performance improvements. The difference between a **Fast Filter** and a **Session Filter** is that **Fast Filters** work at the provider/driver level and are therefore not subject to the Runtime parsing process, which makes them faster, whereas **Session Filters** are applied to an already parsed set of results, which makes them a little slower because of the additional processing time required.

Other **ETW Provider** settings that you can configure for a Live Trace Session are described in the list that follows. Note that the **Provider** tabs of all the **Advanced Settings** dialogs that are referenced in the list items are accessible by clicking the **Configure** link to the right of the providers when they display in the **ETW Providers** list of the **New Session** dialog.

- **System Network** adapter filters and logically ANDed **Fast Filter** group settings — the configuration is accessible from the **Provider** tab of the **Advanced Settings – Microsoft-PEF-NDIS-PacketCapture** dialog for **Local Network Interfaces Trace Scenarios**, as described in the Microsoft-PEF-NDIS-PacketCapture Provider section and in Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

  > **NOTE**
  > The **Microsoft-PEF-NDIS-PacketCapture** provider is available on computers running the Windows 7, Windows 8, or Windows Server 2012 operating system only.

- **Advanced filters** — includes settings for NDIS stack filters; extension layer filters for Hyper-V-Switches that service virtual machines (VMs); and **Direction** (packet traversal), **EtherType**, **IP Protocol Number**, **MAC Address**, and **IP Address** filter settings. The configuration is accessible from the **Provider** tab of the **Advanced Settings – Microsoft-Windows-NDIS-PacketCapture** dialog for the **Local Network Interfaces Trace Scenario**, as described in the Microsoft-Windows-NDIS-PacketCapture Provider section and in Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

- **WFP Layer Set** and **Fast Filter** settings — the configuration is accessible from the **Provider** tab of the **Advanced Settings - Microsoft-Pef-WFP-Message Provider** dialog, as described in the Microsoft-PEF-WFP-MessageProvider section.

- **Hostname** and **Port Filter** settings — the configuration is accessible from the **Provider** tab of the **Advanced Settings – Microsoft-Pef-WebProxy** dialog for the **Pre-Encryption for HTTPS Trace Scenario**, as described in the Microsoft-PEF-WebProxy Provider section.

- **Keyword** event and error **Level** filters — the configuration is accessible from the **ETW Core** tab in the **Advanced Settings** of all provider configuration dialogs; however, not all **ETW Providers** make **Keyword** and **Level** filter settings available, as some providers are not instrumented with them. See System ETW Provider Event Keyword/Level Settings for additional details.

More Information
To learn more about configuring a Live Trace Session, see Capturing Message Data.
To learn more about usage configurations for PEF-based providers and other message providers, see the Built-In Trace Scenarios topic.

Starting a Live Trace Session

After you complete the configuration phase for a Live Trace Session, you can start the session by clicking the **Start** button in the **New Session** dialog, at which time Message Analyzer will begin capturing data. If you have a specific issue that you are trying to resolve, this would be the time to start the function/s or application/s that you suspect are causing a problem.

Note that you can very quickly start capturing data with Message Analyzer by clicking either of the following on the Message Analyzer **Start Page**; however, you cannot set any configuration options for a Live Trace Session when using these methods.

- **Start Local Trace** button — starts a local trace at the Link Layer with the **Microsoft-Windows-NDIS-PacketCapture** provider.

- **Favorite Scenarios** list — starts a local trace with the default **Local Network Interfaces**, **Loopback and Unencrypted IPSEC**, or **Pre-Encryption for HTTPS Trace Scenario** favorites, each of which has a default message provider configuration. Note that you can add other scenarios to the **Favorites** list.

More Information

To learn more details about starting a Live Trace Session, see Performing a Live Capture.

Setting Message Analyzer Global Options

Message Analyzer provides numerous global options that enable you to specify certain default values or make default selections that can affect Message Analyzer performance, display configurations, or feature activation. For example, you can specify a default **Session Viewer**, the default configuration for **Text Log Files**, **Time Display** format, **Decryption** certificate data, **Parsing** options, preview **Features**, **Profiles** to enable, symbol files for parsing WPP-generated events, and so on. You can set these options at any time; however, you would typically do so prior to starting a Live Trace Session or a Data Retrieval Session.

> **IMPORTANT**
>
> If you are enabling preview features on the **Features** tab of the **Options** dialog, as accessible from the global Message Analyzer **Tools** menu, you will need to restart Message Analyzer for the configuration to take effect.

More Information

To learn more about the global Message Analyzer options that you can set, see Setting Message Analyzer Global Options.

Protocol Modules and Specifications

Message Analyzer can display message traffic that is captured from specific protocol modules only if the protocol object model (POM) repository within the PEF architecture contains compiled OPN descriptions representing the architecture, behavior, and data for those protocols. Message Analyzer ships with OPN descriptions for a large number of protocols, such as Microsoft Windows and other common public protocols, in addition to Office, Exchange, SharePoint, and SQL protocols. This enables you to capture a wide array of network protocol and application messages. In addition, to support your data analysis process, Microsoft makes Protocol Technical Specifications available on the Microsoft Developer Network (MSDN) web site, while you can find other standard RFC specifications for public protocols on the Internet.

You can use the technical documents (TDs) provided by Microsoft as references that depict protocol architecture, behavior, and data, as it was designed, to facilitate analysis of the messages you capture with Message Analyzer. For example, you could verify the value of a particular field or confirm the presence of required parameters for a particular method of a specific protocol that is failing to perform properly, although Message Analyzer has a built-in message validation feature that does this automatically.

Integrating Event Tracing

Event tracing functionality is integrated with all message providers that are used by Message Analyzer. Moreover, all Message Analyzer providers are instrumented with ETW technology so that events can be returned in a trace along with network traffic. The Message Analyzer trace model uses ETW to enable integrated capture and display of messages and events from a large number of system components. Whenever you start a Live Trace Session, the underlying message provider/s in the **Trace Scenario** that you select are enabled to an ETW Session Controller, which determines if there are any specific **Keyword** event or error **Level** settings that modify which events are to be returned to the ETW Consumer, which in this case is Message Analyzer. If there are no such settings, then the ETW Session Controller returns all events generated by the component that is instrumented for ETW. Message Analyzer then displays detailed, human-readable information for events at the **ETW** layer that is below the networking stack in all Message Analyzer traces.

At the **ETW** layer in the **Analysis Grid** viewer, ETW messages typically contain an expandable **EventRecord** field and a **Payload** field, the latter of which integrates the network stack. You can see these fields in the **Details Tool Window** if you click an ETW message in the **Analysis Grid** viewer, as shown in the figure that follows.

TIP

If you expand the **EventRecord** node in the **Details** window, you will see the **Header**, which contains fields such as **Size**, **ThreadId**, **ProcessId**, **ProviderId**, and the event **Descriptor**, which contains the fields described in the Event Definition topic of the ETW Framework Conceptual Tutorial.

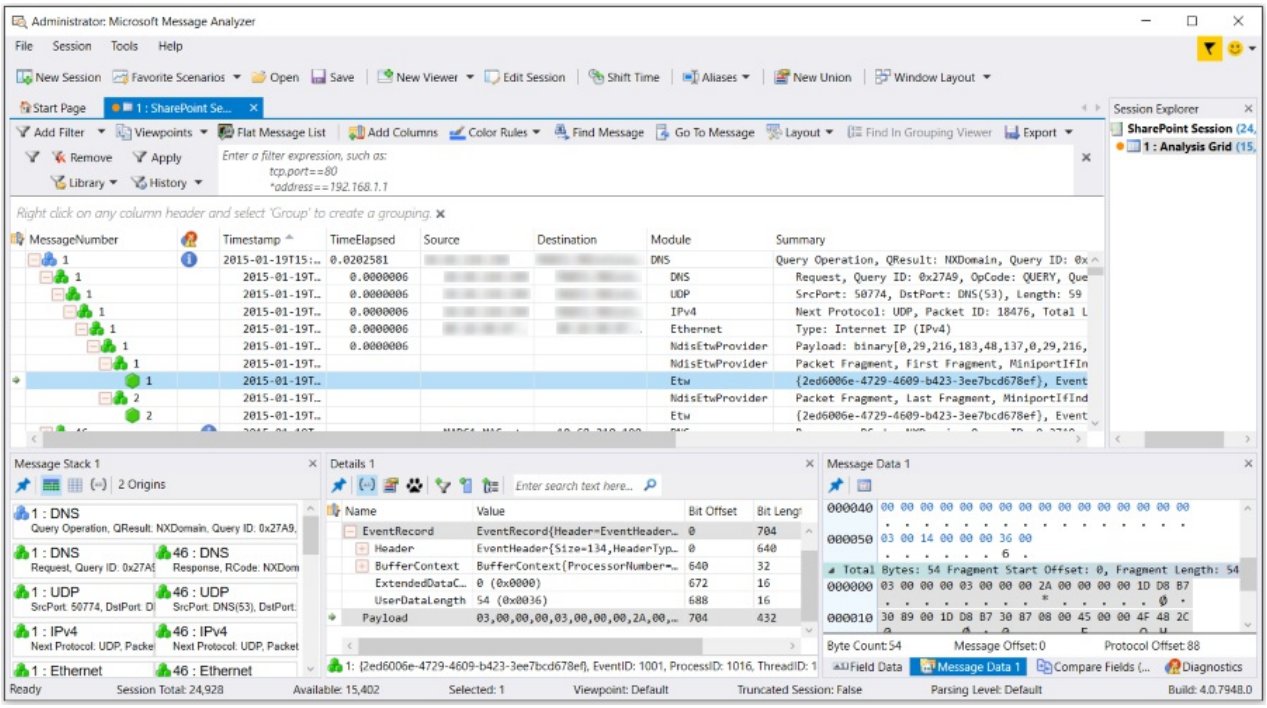

Figure 3: Message Analyzer with Analysis Grid ETW event

For event parsing to be possible, Message Analyzer must generate OPN for any manifest-based system ETW Provider that you employ in a Live Trace Session so that ETW events can be properly parsed by the PEF Runtime. To generate the OPN, manifests for system ETW Providers in use are retrieved so that OPN descriptions can be inferred from them to provide the basis for Message Analyzer to successfully parse event structures. To facilitate this process, the PEF architecture contains an ETW Manifest Import Adapter. This is a protocol object model (POM) adapter that converts an ETW manifest for a given ETW Provider into a POM model, and then publishes it to the PEF Runtime so it can parse and dispatch ETW messages generated by that provider. The OPN actors and endpoints that enable parsing and dispatching messages for an ETW Provider that you specify in a Live Trace Session are dynamically generated at runtime by the ETW Manifest Import Adapter.

TIP

An ETW Provider manifest defines the event descriptions and format in which events are written by the provider. In the current Message Analyzer v1.4 release, you can extend your system with additional system ETW Providers from which Message Analyzer can receive events. If you have a custom ETW Provider that you want to use in a Live Trace Session, you will need to specify a **Guid** and a **Name** for the provider in the **Add Custom Provider** dialog, which displays after you select the **Add Custom Provider** item in the **Add Providers** drop-down list on the **ETW Providers** toolbar of the **New Session** dialog for a **Live Trace**.

However, you might also need to specify a provider manifest so that Message Analyzer can infer an OPN description for the POM to facilitate parsing of the event structure, as described earlier. Message Analyzer will first check to see if the system contains a registered manifest for your provider, and failing that, Message Analyzer looks in the following directory for a manifest:

```
%LocalAppData%\Microsoft\MessageAnalyzer\OPNAndConfiguration\EtwManifests\
```

If Message Analyzer does not find a registered manifest on your system for the custom provider you are specifying, you will need to place the manifest in this directory.

More Information
To learn more about the POM, see the PEF Architecture Tutorial. **To learn more** about ETW, see the ETW Framework Conceptual Tutorial.

Optimizing ETW Session Performance

Message Analyzer also enables you to modify certain aspects of ETW Sessions to focus on capture of specific events and/or to improve performance as follows:

- **ETW Provider** — you can specify the events that you want to receive from a system ETW Provider by configuring **Keyword** and/or **Level** filtering. You can configure **Keyword** and **Level** filters from the **ETW Core** tab in the **Advanced Settings** dialog for the particular message provider that underlies the **Trace Scenario** that you selected, as described in Configuring a Live Trace Session, that is, for system ETW Providers that permit **Keyword** and **Level** filter configuration. Configuring system ETW Provider filtering for event tracing enables you to decrease the event volume and capture time by isolating specific types of events to retrieve in the trace, rather than all events, and enables you to focus your analysis on specific events that you choose.
- **ETW Session Configuration** — you can configure certain aspects of the underlying ETW Session in which an ETW Provider participates to enhance session performance. This mainly involves adjusting settings for the ETW buffer configuration of the ETW Session that is managed by an ETW Session Controller. These adjustments are available from the Message Analyzer **ETW Session - Advanced Configuration** dialog that is accessible by clicking the **Configure ETW Session** button on the **ETW Providers** toolbar in the **New Session** dialog, as shown in Figure 2.

More Information
To learn more about optimizing an ETW Session, see Specifying Advanced ETW Session Configuration Settings.
To learn more about how system ETW Providers function in the ETW framework, see the ETW Framework Conceptual Tutorial.
To learn more about configuring system ETW Providers, including **Keyword** and **Level** filters, see Adding a System ETW Provider and System ETW Provider Event Keyword/Level Settings.

Using MOF-Based ETW Providers

Message Analyzer also supports registered event providers on your system that use the managed object format (MOF) schema as the basis of generating their events. Event providers that use the MOF schema are typically employed in systems that are managed by Windows Management Instrumentation (WMI). These providers appear in the **Add System Providers** dialog along with various other types of providers, such as those that are manifest-based. The **Add System Providers** dialog displays after you click the **Add Providers** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog and you select the **Add System Providers** item. Because Message Analyzer supports MOF schema, events that are captured by Message Analyzer from MOF-instrumented providers can be fully parsed. Without MOF support, messages that are captured from MOF-based providers would be displayed as simple ETW messages with a summary string and no additional parsing of event fields.

To provide support for MOF-instrumented providers, including fully parsing events from such providers, Message Analyzer uses an extension to the existing ETW adapter. This adapter normally handles ETW providers that have a manifest that is created at the time the provider is instrumented for ETW. When an ETW event arrives, Message Analyzer checks to see whether an OPN description exists that can parse the event. If an OPN description cannot be found, then Message Analyzer attempts to retrieve the manifest-based event schema, from which it can generate OPN. In a similar manner, Message Analyzer does the following to support MOF when events arrive:

- Verifies whether events are generated by an MOF-based provider.
- Checks the local system for an existing OPN description that can parse the events.

- Uses the extended version of the ETW adapter to generate an OPN description based on the MOF schema of the provider, if an existing OPN description was not found.

Detecting MOF Schema

In Message Analyzer, there are typically three sources from which MOF events can derive, including live traces, saved trace files such as the native Message Analyzer parsed format (.matp), and saved trace files in other supported formats such as .matu, .etl, and .cap. As previously indicated, if there is an existing OPN module (see Protocol Modules and Specifications) that can consume the events, then the events are parsed according to the OPN description and background generation of OPN is not required. However, if there is no existing OPN module to parse the events, Message Analyzer then attempts to locate the MOF schema as follows:

- **Live trace** — when you run a Live Trace Session that utilizes MOF-based event providers, the locally installed MOF schemas are retrieved from the appropriate event provider/s that are installed on the local machine, and OPN descriptions for the provider events are automatically generated for parsing the event fields.
- **Saved .matp files** — if one or more MOF schemas were used to parse messages from an MOF provider when a trace is taken with Message Analyzer, the schemas become part of the .matp trace file when it is saved in the same format. The schema is thereafter provided to Message Analyzer at the time the .matp trace file is loaded, making it independently available to facilitate event parsing whether or not MOF schemas exist on the local system or were deployed during Message Analyzer installation.
- **Saved non-.matp files** — these files will not contain the embedded schema information, therefore Message Analyzer looks up local files deployed during installation. If a local .mof file is discovered, it is used as the MOF schema from which an OPN description is generated for parsing events. Otherwise, the system MOF schema is retrieved and used in a similar manner.

> **NOTE**
> If Message Analyzer requires a MOF schema for a provider that is installed on the local system and cannot find one, then Message Analyzer will display simple ETW messages only, with minimal parsing for that provider's messages.

Deploying a Custom MOF Provider

If you have a custom MOF-based provider that you want to deploy on your local system, you can use the WMI compiler tool *mofcomp.exe* to register your provider and its MOF schema. Thereafter, Message Analyzer will be able to locate the MOF schema, should an OPN description need to be created to parse the MOF-based events of the provider. You will find the *mofcomp.exe* tool in the following directory on your computer:

```
C:\Windows\System32\wbem\
```

More Information

To learn more using the *mofcomp.exe* tool, see mofcomp in the *WMI Command Line Tools* topic on MSDN.

Retrieve Message Data

This section briefly describes how to create a Data Retrieval Session, how to create a message collection from a set of specified input files (or by selecting a subset of specified input files), the features you can use to *select* specific data from a collection of messages in one or more input files, in addition to how to parse text-based log files (with a .log extension). The subject matter is discussed in the following topics.

Loading Data into Message Analyzer
Acquiring Input From Other Data Sources
Selecting Data to Retrieve
Parsing Input Text Log Files

Loading Data into Message Analyzer

When you start a Data Retrieval Session, the configuration of which is shown in the figure that follows, you can load data from saved trace files and logs into Message Analyzer, which includes .matu, .matp, .etl, .cap, .pcap, .log files, and others, as described by the table in Locating Supported Input Data File Types. After clicking the **Add Files** button on the **Files** tab in the **New Session** dialog for a Data Retrieval Session, you can navigate to target files that contain the data you want to load into Message Analyzer. After the files containing the target data display on the **Files** tab, you can also specify subsets of those files in your **Files** list to create message collections that target specific data to be loaded into Message Analyzer and parsed. To create a subset, you simply select the check box to the left of the file that contains the data you want to load. Note that a Data Retrieval Session enables you to aggregate and merge message data from multiple data sources that include various types of log files and traces.

To create a uniform analysis context for your data, you can apply a common filtering configuration to each collection of input files that you specify as a separate **Data Source** in the **New Session** dialog. Such filtering includes specifying a **Session Filter** or **Parsing Level**. However, you have the option to apply a *different* **Time Filter** configuration to each **Data Source** (on a **Files** tab), which gives you the flexibility to aggregate messages from multiple data sources in a specific window of time. You can also specify a data viewer of choice that applies to all **Data Sources**, by choosing it from the **Start With** drop-down list in the **New Session** dialog.

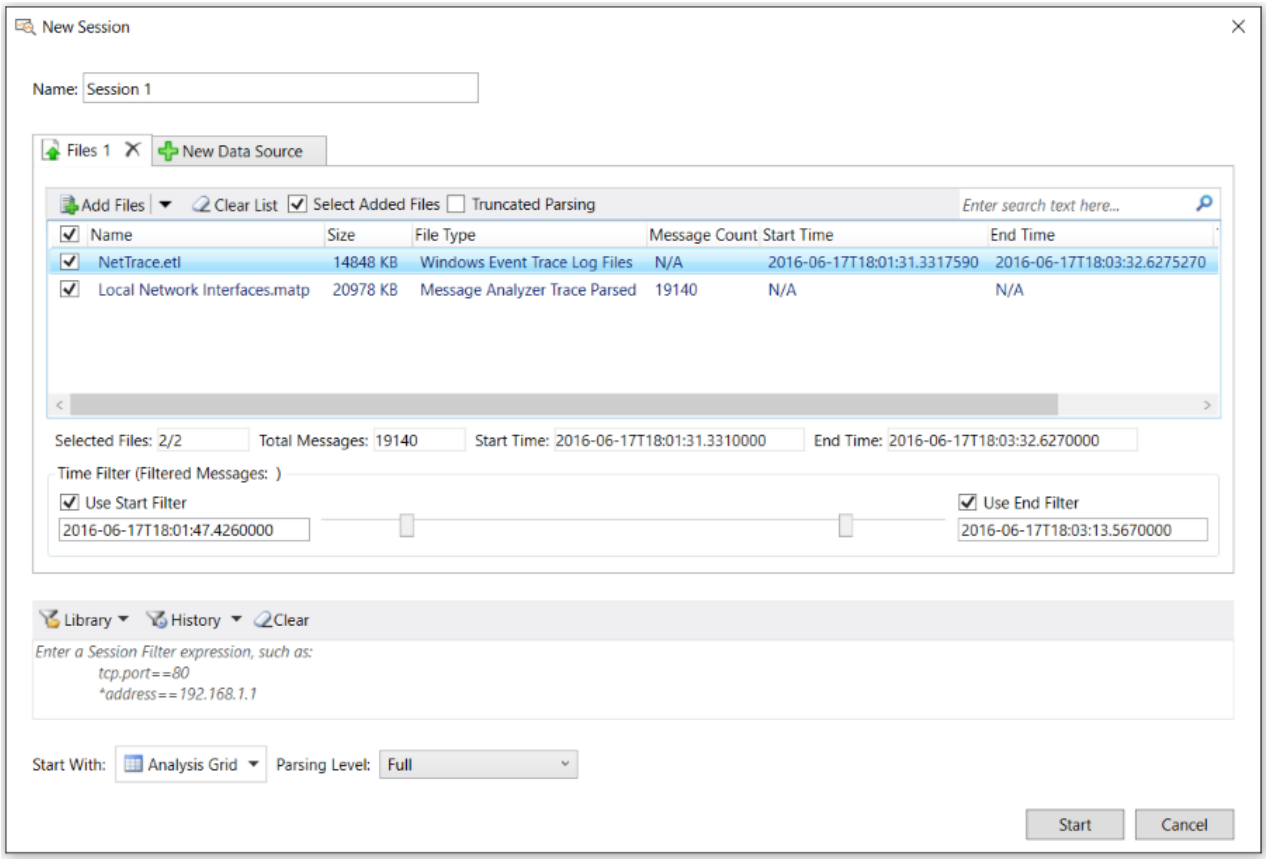

Figure 4: Message Analyzer Data Retrieval Session configuration

More Information

To learn more about working with a Data Retrieval Session, see Configuring a Data Retrieval Session.

Acquiring Input From Other Data Sources

Message Analyzer can load and process data from other input **Data Source** types besides trace files and common log files. The other sources with which Message Analyzer can work include the following:

- **Azure Tables** — Message Analyzer enables you to load input data from Azure tables. You can do this by creating an Azure input configuration from the **New Session** dialog that specifies Azure **Account**

connection information and a **Table Name**.

Azure Storage Blobs — Message Analyzer enables you to browse for, select, and view data from log files that are stored in Azure binary large object (BLOB) containers. You can do this my making use of the **File Selector** dialog, which is accessible by clicking the **From Other File Sources** item in the **Open** drop-down list that appears on the global Message Analyzer **File** menu.

More Information
To learn more about working with Azure data as an input source to Message Analyzer, see Handling Azure Data.

- **Event Logs** — Message Analyzer enables you to load system event data that is typically displayed in the Microsoft Event Viewer.

More Information
To learn more about working with system event data in Message Analyzer, see [Loading System Event Log Data] (loading-system-event-log-data.md).

- **PowerShell** — enables Message Analyzer to acquire input data through PowerShell. For example, you can import data with a saved PowerShell script file that you target as a supported input file type (*.ps1) through a Data Retrieval Session. The script contained in such a file may invoke specific processes or functions which return data that you can view in Message Analyzer. You also have the option to use a PowerShell interface that is built into Message Analyzer to create a PowerShell query that returns its data to a viewer such as the **Analysis Grid**.

More Information
To learn more about working with PowerShell as an input source to Message Analyzer, see [Deriving Input Data with PowerShell Scripts] (deriving-input-data-with-powershell-scripts.md).

- **SQL** — enables you to retrieve data from a SQL database table by using a built-in interface to provide connection information, SQL query code, and a reference timestamp.

More Information
To learn more about working with SQL data as an input source to Message Analyzer, see Loading SQL Data.

- **OMS** — enables you to load data from Operations Management Suite (OMS) logs through a search interface to OMS Log Analytics that Message Analyzer provides.

More Information
To learn more about working with OMS data as an input source to Message Analyzer, see Loading OMS Log Data.

- **WPP-generated events** — enables you to parse and display Windows software trace preprocessor (WPP)-generated events in Message Analyzer, which can capture these events live or load them from a saved event trace log (ETL) file. However, you will need to provide a program data base (PDB) or trace message format (TMF) file that defines the event structure and format so that Message Analyzer can parse the WPP-generated events.

More Information
To learn more about the configuration required for Message Analyzer to parse WPP-generated events, see [Loading WPP-Generated Events] (loading-wpp-generated-events.md).

Selecting Data to Retrieve

You can also *select* specific data to retrieve from a target message collection while blocking all other messages that do not meet the filtering criteria that you define, by using a **Session Filter**, **Time Filter**, or a **Parsing Level**. A **Session Filter** narrows the scope of data retrieval to only the message types that meet the criteria of a Filter that you manually define, or one that you select from the centralized filter **Library** in the lower section of the **New Session** dialog. A **Time Filter** enables you to specify a window of time in which to view data in a correlated target message collection that can consist of one or more sources from which you load data into Message Analyzer. A **Parsing Level** enables you so specify how far up the network stack that Message Analyzer will parse, which creates a focused set of messages that temporarily eliminates all other messages above the specified **Parsing Level**. For example, you might set the **Parsing Level** known as **Network Analysis** to create a set of results that focuses on the Network and Transport Layer messages.

More Information

To learn more about configuring a Data Retrieval Session, see Retrieving Message Data.

To learn more about how to use a **Session Filter** in a Data Retrieval Session, see Applying a Session Filter to a Data Retrieval Session.

To learn more about how to use a **Time Filter** in a Data Retrieval Session, see Applying an Input Time Filter to a Data Retrieval Session.

To learn more about how to work with **Parsing Levels**, see Setting the Session Parsing Level.

Parsing Input Text Log Files

If you have a text-based log file containing log entries that you want to view, Message Analyzer enables you to load and view the data from the log file, but you will need to specify an OPN configuration file to drive the process. Message Analyzer provides several built-in configuration file types that you can select from in the **Text Log Configuration** column that appears below the toolbar on the **Files** tab of the **New Session** dialog for a Data Retrieval Session, as shown in the figure that follows.

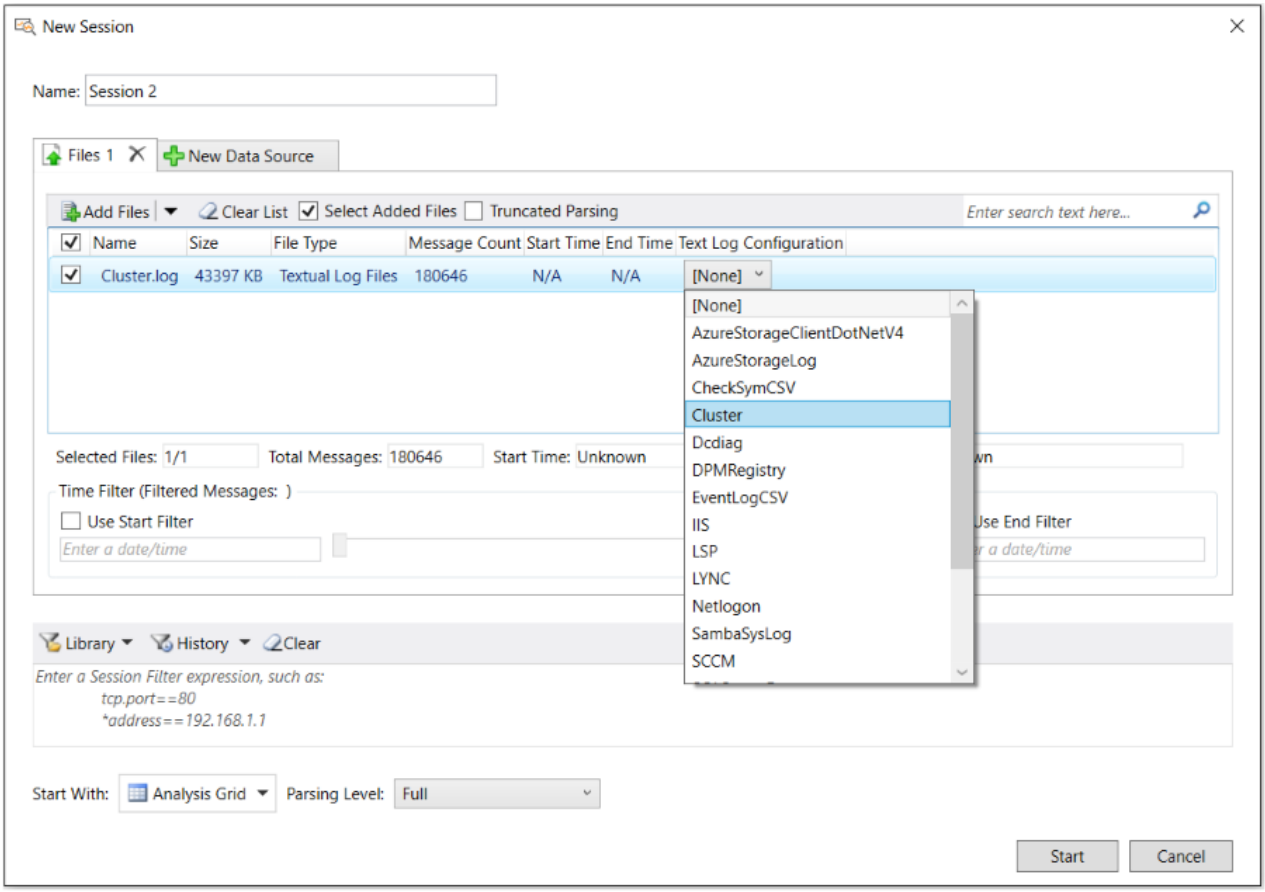

Figure 5: Message Analyzer Textlog parsers

The drop-down list shown in the figure is populated with common built-in configuration files that are available for selection only after you click the **Add Files** button and retrieve a *.log file that contains the data you want to load

into Message Analyzer. The built-in configuration files are described in the subsection "Built-In OPN Configuration Files" that immediately follows; however, if you have a text-based log file that contains log entries in a unique/proprietary format, it is likely that you will need to create a custom OPN configuration file so that Message Analyzer can parse your log, as described in Opening Text Log Files.

Built-In OPN Configuration Files The built-in OPN configuration file types that are currently available for selection are specified in the list that follows. A short description of the purpose of each configuration file type is included:

- **AzureStorageClientDotNetV4** — provides the OPN configuration that parses Azure .net client storage logs.
- **AzureStorageLog** — provides the OPN configuration that parses Azure .log files that are saved in BLOB containers.
- **CheckSymCSV** — provides the OPN configuration that parses the CSV output of the Exchange CHECKSYM utility, which is commonly used to perform file version and checksum comparisons of binaries and configuration files.
- **Cluster** — provides the OPN configuration that parses Cluster text logs.
- **Dcdiag** — provides the OPN configuration that parses the output of the Domain Controller Diagnostics Tool (Dcdiag).
- **DPMRegistry** — provides the OPN configuration that parses special registry output text logs for the Data Protection Manager (DPM) component.
- **EventLogCSV** — provides the OPN configuration that parses traces that are exported as a CSV file, but with more value than a regular CSV file.
- **IIS** — provides the OPN configuration that parses text logs generated by IIS web servers.
- **LSP** — provides the OPN configuration that parses text logs generated by the Local Security Authority (LSA) component, which applications can use to authenticate and log users on to the local system. The log files provide access to some data in clear text that is otherwise encrypted by messages on the wire. Administrative privileges are required to view these logs.
- **LYNC** — provides the OPN configuration that parses UCCAPI logs from the Lync client application.
- **Netlogon** — provides the OPN configuration that parses Netlogon logs for diagnosing logon issues on domain controllers.
- **SambaSysLog** — provides the OPN configuration that parses SambaSysLog text logs generated for Unix and Linux machines.
- **SCCM** — provides the OPN configuration that parses System Center logs.
- **SQLServerError** — provides the OPN configuration that parses SQL Server error logs.
- **SQLServerSetup** — provides the OPN configuration that parses SQL Server setup logs.
- **ULS** — provides the OPN configuration that parses SharePoint logs.
- **VMM** — provides the OPN configuration that parses System Center Virtual Machine Manager logs.
- **DefaultSimpleLogFileReaderConfig** — a generic configuration file that can parse most text logs, for example those from a domain controller.

NOTE

With exception of the configuration files for Azure storage logs, the listed text log configuration files are contained in the Message Analyzer **Device and Log File Version 1.4** asset collection that you can configure for automatic downloads and updates from a Microsoft web service through the Sharing Infrastructure. The configuration files for Azuze storage logs are contained in the **Azure Storage Parsers Version 1.0** asset collection. The management features for the Azure storage parsers and all other Message Analyzer asset collections are available from the **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu.

More Information

To learn more about managing Message Analyzer asset collections, including downloading and auto-syncing any collection for automatic updates, see Managing Message Analyzer Assets.

Selecting Versus Creating an OPN Configuration File

The built-in text log OPN configuration files are named in such a way that it should be obvious which one to select for your text log. For example, a Cluster text log will use the **Cluster** configuration file, the IIS text log will use the **IIS** configuration file, and so on.

If none of the built-in text log configuration files apply to your text log, then you can create a new one that is specifically designed to parse the data in your text log, as described in Opening Text Log Files. Whenever you create a new configuration file for a text log, it is added as an item to the **Text Log Configuration** drop-down list that appears below the toolbar on the **Files** tab of the **New Session** dialog. It is also added to the **Default text log configuration** drop-down list in the **Text Log Files** pane on the **General** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. From the latter drop-down list, you have the option to set a specific configuration file as the global default for all text log files from which you will load data into Message Analyzer. This makes it convenient if you work with a particular type of .log file consistently.

OPN Configuration File Contents

A configuration file contains a description of the log's messages in OPN and RegEx notation, which ensures that text log data that is loaded into the system can be properly parsed and then displayed in Message Analyzer. The text-based log data is loaded into the Message Analyzer Runtime through a Log File Adapter and the OPN configuration file drives the process. The message definitions contained in the OPN configuration file are compiled by the OPN Compiler to confirm the validity of the configuration file and the integrity of the OPN description that will reside in the POM, which is referenced by the Runtime when the parsing process begins for your text log.

To create an OPN configuration file, you will need to identify each unique log entry and map it to a message structure. You can do this with RegEx notation, which is designed for matching strings of text. RegEx provides the functionality you will need to match data through the mechanism of capture variables, which you can use to map extracted log file data to field names that you define in OPN; in turn, these become data columns in the **Analysis Grid** viewer.

TIP

Message Analyzer also supports loading regular comma-separated-value (CSV) and tab-separated-value (TSV) data file formats directly, without the need for an OPN configuration file.

More Information

To learn more about how to create an OPN configuration file, download the [OPN Configuration Guide for Text Log Adapter](https://download.microsoft. com/download/C/D/E/CDED67DB-2C74-4FE4-B184-123CEE0E273F/OPN%20Configuration%20Guide%20for%20Text%20Log%20Adapter%20V2.docx) document.
To learn more about other OPN configuration file requirements, see the Addendum 1: Configuration Requirements for Parsing Custom Text Logs topic.

Edit Message Data

Message Analyzer enables you to edit the data of any Live Trace Session or Data Retrieval Session. You can achieve this by modifying the session configuration and applying the changes you make. To modify the configuration for either of these types of sessions, simply click the **Edit Session** button on the global Message Analyzer toolbar to display the **Edit Session** dialog. The session configuration that displays in the **Edit Session** dialog depends on the session viewer tab that has focus (viewer tabs are below the global Message Analyzer toolbar). Note that only one session configuration exists for a specified session, regardless of how many data viewers are open in that session.

Note that the **Edit Session** dialog for a Live Trace Session is similar to the **New Session** dialog shown earlier in Figure 2, while the **Edit Session** dialog for a Data Retrieval Session is similar to the **New Session** dialog shown earlier in Figure 4, with exception of the **Restricted Edit** information bar.

Editing a Data Retrieval Session
When you open the **Edit Session** dialog for a Data Retrieval Session, it opens in **Restricted Edit** mode, which means you can add more files to the files list and display the data contained in such files without incurring a full reload of data. When the session changes take effect (after you click **Apply** in the dialog), the data from the new input files is appended to the existing data file import results. However, if you click the **Full Edit** button on the information toolbar, you have additional options to modify the session configuration.

Similar to initial configuration of a Data Retrieval Session, the changes you can make to the Data Retrieval Session configuration include not only more input data files, but one or more of the following as well:

- **Time Filter** — configure a window of time in which to view messages.
- **Session Filter** — apply a built-in or custom-written Filter to the session configuration.
- **Parsing Level** — apply a parsing level to the session to establish the level up to which Message Analyzer will parse messages. Accessible from the **Parsing Level** drop-down list in the **New Session** dialog.
- **Truncated Parsing** — specify the **Truncated Parsing** option for input files that contain truncated messages. Results in improved performance.

When you edit a Data Retrieval Session with any of these features, the session data will be reloaded with the specified asset/s applied, for example, a **Session Filter** and/or a **Time Filter**. This enables you to modify the session results and obtain a different view of the data. You can edit a Data Retrieval Session in the specified manner as many times as you wish.

Editing a Live Trace Session
When you open the **Edit Session** dialog for a running, paused, or stopped Live Trace Session, it opens with no editing restrictions; this means you can make modifications to the session configuration and **Apply** them as required. Whatever changes you make to a *running* Live Trace Session will take effect on subsequent messages that the Live Trace Session in progress is capturing, that is, after you click **Apply** in the **Edit Session** dialog. If you edit a *stopped* or *paused* Live Trace Session, the changes do not take effect until you restart the session, either by clicking the **Restart** button or the **Pause/Resume** button, respectively, on the global Message Analyzer toolbar.

Similar to initial configuration of a Live Trace Session, the changes that you can make to the Live Trace Session configuration include one or more of the following:

- **Trace Scenario** — replace any existing **Trace Scenario** with a new one, as you can only have a single scenario in the session configuration at any one time.
- **Specific ETW Provider/s** — add system ETW Providers or custom ETW Providers that you specify from the **Add Providers** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog.
- **Advanced Settings** — specify advanced settings for a selected provider. Accessible from the **Configure**

link next to each message provider in the **ETW Providers** list. Displays the **Advanced Settings** dialog, from where you can specify ETW **Keyword** or **Level** filters and other provider filtering settings that are unique to each provider type.

- **ETW Session Configuration** — specify settings for the ETW session such as buffer size, buffer count, and buffer flush timing for ETW events. Accessible from the **ETW Session - Advanced Configuration** dialog that displays when you click the **Configure ETW Session** button on the **ETW Providers** toolbar.
- **Session Filter** — apply a built-in or custom-written Filter to the session configuration.
- **Parsing Level** — apply a parsing level to the session that establishes the level up to which Message Analyzer will parse messages. Accessible from the **Parsing Level** drop-down list in the **New Session** dialog.
- **Target Computers** — specify one or more target computers on which you want to capture data. Accessible from the **Edit Target Computers** dialog which appears after you click the **Edit** button next to the **Target Computers** text box in the **New Session** dialog.
- **Data Source** — add and configure a new **Data Source** through which you can capture messages. Accessible from the **New Data Source** tab in the **New Session** dialog.

More Information
To learn more about editing a session, see Editing Existing Sessions.
To learn more about **Trace Scenarios**, see the Built-In Trace Scenarios topic.
To learn more about working with system ETW providers, see Adding a System ETW Provider.
To learn more about **Advanced Settings** for system ETW Providers, see the topics Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog or Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.
To learn more about ETW session configuration, see Specifying Advanced ETW Session Configuration Settings.
To learn more about **Session Filters**, see Working with Session Filters in a Live Trace Session and Applying a Session Filter to a Data Retrieval Session.
To learn more about **Time Filters**, see Applying an Input Time Filter to a Data Retrieval Session.
To learn more about **Parsing Levels**, see Setting the Session Parsing Level. **To learn more** about how to configure a session for capturing traffic on remote computers, see Configuring a Remote Capture.
To learn more about how to configure multiple data sources, see Configuring Session Scenarios with Selected Data Sources.

View Message Data

Message Analyzer provides many different data viewers and **Layouts** in which to present data that you capture from a Live Trace Session or load from a Data Retrieval Session. It even provides **Window Layout** presets that each have the **Analysis Grid** viewer in common along with varying arrangements of Message Analyzer **Tool Windows**, to create custom working environments that suit the type of troubleshooting and analysis that you typically perform.

The **Analysis Grid** viewer is the main analysis surface provided by Message Analyzer. It has a tree grid configuration with selectable column **Layouts** that expose a wide array of data fields that are useful for analyzing different types of data. It also uses a unique message encapsulation and stacking scheme to organize data so that important information is readily accessible at the top-level of the display. This configuration eliminates the time normally needed to search for related but dispersed information in trace files that contain a high volume of messages.

For example, in the **Analysis Grid** viewer, the message stack is encapsulated under expandable top-level transactional messages and Operations, while message fragments at the Transport Layer are reassembled as part of the PEF Runtime parsing process. Moreover, each line of data in the **Analysis Grid** viewer displays message

data as expandable top-level parent nodes that contain all the child node (message stack) messages and message fragments that were involved in a particular transaction or Operation. Note that you do have the option to display message data in any Message Analyzer data viewer that you select, however, the encapsulation and stacking scheme only exists in the **Analysis Grid** viewer. The **Analysis Grid** viewer, which is shown in the subsection that follows, is fully described in the Analysis Grid Viewer topic.

The material of this section is covered in the following topics.

Organizing Messages in the Analysis Grid Viewer
Grouping Messages in the Analysis Grid Viewer
Grouping Messages in the Grouping Viewer
Applying Viewpoints
Viewing Message Details
Viewing Other Message Data
Viewing Data from Multiple Sessions
Limiting the Scope of Applied Assets
Driving Interaction Between Data Viewers
Using Window Layouts
Using Message Analyzer Profiles

Organizing Messages in the Analysis Grid Viewer

As indicated earlier in this Tutorial, the overarching approach to analysis in Message Analyzer is to bring key data into focus wherever possible to make it more accessible, which streamlines and therefore expedites the analysis process. In keeping with this approach, the Message Analyzer Runtime creates Operation nodes for protocols that use the request/response conversation architecture, such as DNS, HTTP, SMB2, and so on. The Runtime also reassembles the message stack, including fragments (such as TCP virtual segments), and by default hides them under expandable top-level message nodes in the **Analysis Grid** viewer for easy access. For example, a top-level message node could be an Operation that encapsulates a request/response message pair, under each of which resides the message stack and fragments that supported the Operation.

By organizing messages this way, you can easily determine such important values as the **ResponseTime**, which can tell you how long it is taking to receive the first server response to a request message; by utilizing this feature, you can avoid searching through potentially hundreds, if not thousands of messages to find such a response message. The **ResponseTime** is important to analysis because it can indicate how long a service is taking to respond, which can rule out network issues while potentially indicating server issues instead. Note that by sorting the **ResponseTime** column in the **Analysis Grid** viewer, you can readily determine the specific Operations that had the longest response times. However, to view **ResponseTime** data, you will need to add this **Global Annotation** from the **Field Chooser Tool Window** as a new **Analysis Grid** column.

Another important value is the **ElapsedTime**, which can tell you how long an Operation is taking to complete; this includes how long it took to receive all the associated message fragments. If the **ElapsedTime** is a comparatively high value with respect to **ResponseTime**, this could be an indication of a network issue. Also, by performing a sort of the **ElapsedTime** column in the **Analysis Grid** viewer, you can determine the specific Operations (with fragments) that took the longest to complete — which can be a cue for further investigation.

More Information

To learn more about the **Field Chooser**, see Using the Field Chooser and the Field Chooser Tool Window topics.
To learn more about **ResponseTime**, see the Average Response Time for Operations topic.

Another important aspect of organizing messages as top-level nodes with an encapsulated network stack, is that it enables you to have immediate access to stack messages for quick analysis of details, whereas in other tools such as Microsoft Network Monitor, stack messages are typically chronologically dispersed across a set of trace results, making them difficult to find and correlate to a top-level transaction.

The top-level node and encapsulation configuration also provides a visual cue of **Diagnosis** messages that Message Analyzer drives to top-level, even when those errors occur at deeper stack levels. Although such errors can be initially hidden in the child messages that make up the encapsulated stack, nevertheless you can simply click the error icon at top-level to display the **Diagnosis** tab of the inline message details where you can review the error text, correlated to the Operation or other top-level message at hand. This provides easy access to error information without having to search through a multitude of messages to discover it. If you want to know exactly where in the stack the error occurred, you can merely expand the child message nodes until you find the specific message that contains the error icon that initially displayed at top-level.

As an example of the benefits of the described message organization, the figure that follows shows the **Analysis Grid** viewer with an expanded SMB2 Operation node containing a request/response message pair and the expanded message stack showing message fragments for a response message (which also has a Diagnosis error).

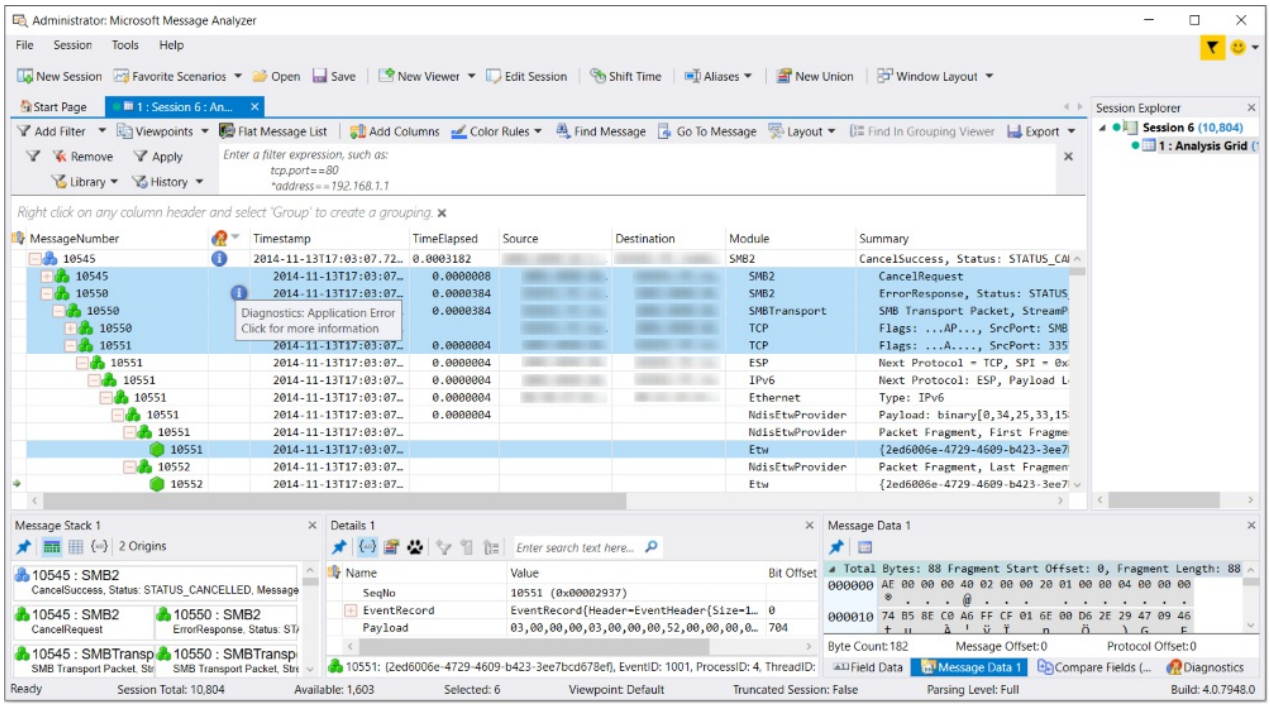

Figure 6: Message Analyzer expanded Operation node, message stack, fragments, and Diagnosis error

In this figure, note the highlighted **SMB2** request and response message numbers **10545** and **10550**, respectively, that are encapsulated under a top-level **SMB2** Operation node which is designated by a blue-cubed icon. Also, the response message stack is expanded to show the stack messages which includes the message fragments that consist of one **SMBTransport** and two **TCP** fragments. Also note that the Operation message number **10545** contains a blue **Diagnosis** error icon that Message Analyzer bubbled up from the response message **10550** where the error actually occurred, so that you can see it at-a-glance from the top-level Operation node, even when this node is in the unexpanded state. Note that you can click the **Diagnosis** error icon in either location for more information, which in this case happens to specify an **Application** error with a **Warning** level.

Grouping Messages in the Analysis Grid Viewer

Another feature that is important to data analysis is the **Group** feature. By right-clicking selected **Analysis Grid** viewer columns in succession and selecting the context menu **Group** item for each one, you can create a data display of nested groups that provides a convenient way to organize and explore targeted trace data. As an additional example of grouping, you could create IPv4 **Network** and TCP **Transport** groups in the **Analysis Grid** viewer by executing the **Add as Grouping** command from the right-click context menu for these fields in the **Field Chooser** window to **Group** your data into these field-categories. This quickly organizes your data into groups of IP conversations that took place across a trace, with the TCP ports that supported those conversations nested within each IP group, resulting in a unique analysis perspective. However, to perform this operation, the **Analysis Grid** viewer must be in focus.

The **Analysis Grid** viewer **Group** feature essentially categorizes your data according to the field data you are

grouping and the order in which you group it. The **Group** feature enables you to extract all the data from your trace into the categories that you establish through the grouping process, which results in bringing hidden or dispersed trace messages into what you might call a "categorical focus". A figure that shows the results of grouping data in the **Analysis Grid** viewer is provided in the topic that is referenced in the "More Information" section that immediately follows.

More Information

To learn more about the **Group** function in the **Analysis Grid** viewer, see Using the Analysis Grid Group Feature.

Grouping Messages in the Grouping Viewer

You can also make use of the **Grouping** viewer, which has a set of built-in view **Layouts** that render your message data into a separate view of predefined nested Group configurations that integrate and interact with other data viewers to create unique analysis contexts. You can also create and save your own **Grouping** view **Layouts** that you customize to your environment based on message fields that you select from the **Field Chooser** window. The **Grouping** viewer is accessible from the **New Viewer** drop-down list that appears on the global Message Analyzer toolbar. A **Layouts** drop-down list is included on the **Grouping** viewer toolbar.

The **Grouping** viewer has functional similarities with the **Analysis Grid** viewer **Group** feature, in that they both enable you to create nested groups of data that are hierarchically categorized by the message fields that you use for the groupings. With the **Grouping** viewer, you can organize your traffic into summary hierarchies based on built-in or custom-designed **Grouping** view **Layouts** that are configured with message field groups in nested configurations. You can also manually adjust (pivot) your group **Layout** by dragging and dropping Group labels to change the nesting order and obtain different message correlation configurations that result in unique analysis contexts.

Grouping Viewer Advantages The following summarizes the advantages of viewing data with the **Grouping** viewer, where you can:

- Organize data into unique hierarchies to expose targeted information that you can quickly extract from large data sets, which can otherwise be difficult to do.
- Arrange nested Group configurations so you can isolate messages of interest into specified categories, which enable you to drill down into the nested Groups to obtain a concise analytical focus.
- Create a coherent analysis context for messages that can otherwise appear as disassociated.
- Locate the Group(s) with the highest message volumes for performance assessments.
- Correlate message volumes in different Groups.
- Drive the display of Group messages into the **Analysis Grid** viewer for assessment of message details.

> **NOTE**
>
> Every Message Analyzer installation provides a default **Message Analyzer Grouping View Layouts** asset collection that appears in the **Asset Manager** dialog, where you can manage downloads and the auto-sync feature to update the collection. You can access the **Asset Manager** dialog from the global Message Analyzer **Tools** menu.

Example Layout

The figure that follows shows an example of the **Grouping** viewer with the **File Sharing SMB/SMB2** view **Layout**, that displays three nested Groups as identified by the labels below the **Grouping** viewer toolbar: **SessionIdName**, **TreeIdName**, and **FileName**. This built-in **Grouping** viewer **Layout** was pre-configured with these Group names by locating the corresponding fields in the **Field Chooser** window under the **QueryDirectoryRequest** node of the **SMB2** message hierarchy. You can add other related SMB2 fields as Groups at your discretion, by right-clicking a particular field in **Field Chooser** and selecting the **Add As Grouping** item in the context menu that appears. This action will create a new nested Group identified by the field

name that you selected, at which time, the **Grouping** viewer data display will be refreshed to include the new Group.

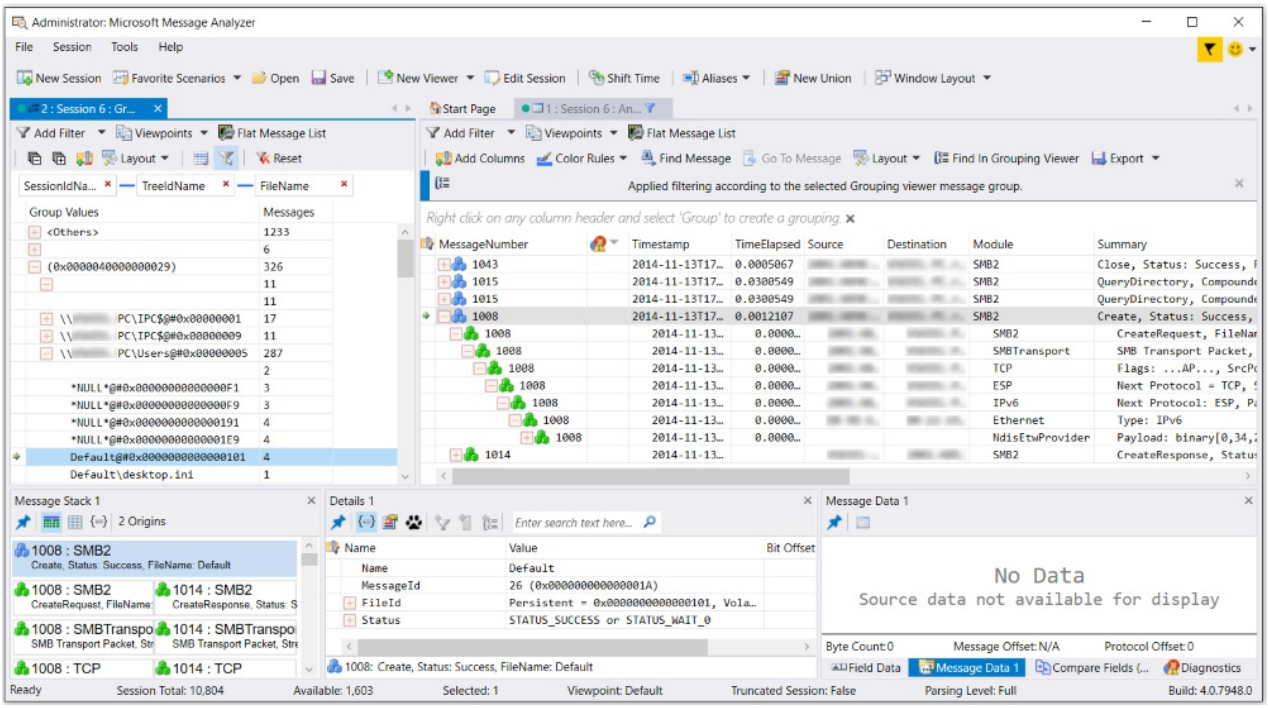

Figure 7: Message Analyzer Grouping viewer selection interactively driving Analysis Grid message display

In this figure, the **Grouping** viewer shows a file name selected in the **FileName** group, which is ***NULL*@#0x0000000000000191**, and this group is nested under the **TreeIdName** Group value of **\\PC\Users@#0x00000005**, which in turn is nested under the **SessionIdName** Group value of **(0x0000040000000029)**. The number of messages associated with this particular SMB2 operation is specified in the corresponding row under the **Messages** column of the **Grouping** viewer. Whenever you select a row of data in any Group in the **Grouping** viewer, the corresponding messages are interactively displayed in the **Analysis Grid** viewer for further analysis of message details. message stack, message data, field data, diagnostics, and so on.

More Information

To learn more about the **Grouping** viewer, see the Grouping Viewer topic.

Applying Viewpoints

To simplify troubleshooting, Message Analyzer provides the **Viewpoints** feature that enables you to examine network traffic from the perspective of a protocol. An applied **Viewpoint** enables you to bring the messages of a particular protocol or module into focus for targeted analysis. By applying a built-in **Viewpoint** from the **Viewpoints** drop-down list on the Filtering toolbar shown in the figure below, you can focus on specific messages at top-level in the **Analysis Grid** viewer with no layers above them, as defined by the applied **Viewpoint**. Moreover, because the **Viewpoint** temporarily removes all messages above the applied protocol **Viewpoint**, only the protocol messages associated with the applied **Viewpoint** appear at top-level in the **Analysis Grid** viewer. This feature is advantageous when you have higher-layer traffic that obscures the underlying messages that you want to troubleshoot. For example, if you were interested in focusing on SMB messages at the Application Layer, you could apply the **SMB/SMB2 Viewpoint** as shown in the figure that follows. Upper sublayer protocols such as RPC will be removed from the display, as you will see SMB messages only.

> **NOTE**
>
> Every Message Analyzer installation provides a built-in **Message Analyzer Viewpoints** asset collection that appears in the **Asset Manager** dialog, where you can manage downloads and the auto-sync feature to update the collection.

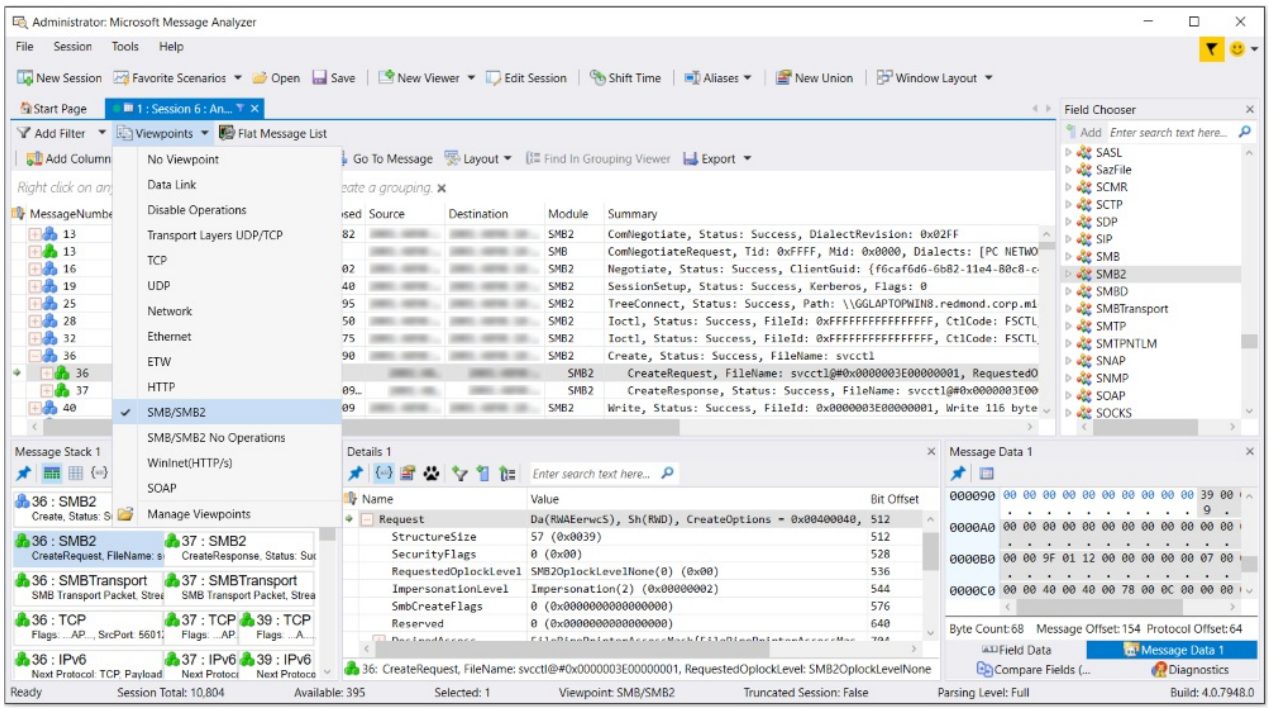

Figure 8: Message Analyzer SMB/SMB2 Viewpoint applied

In this figure, you can see that only **SMB2** messages display in the **Analysis Grid** viewer when the **SMB/SMB2 Viewpoint** is applied, as indicated by the check mark in the **Viewpoints** drop-down list. To return to your original message set, simply select the **No Viewpoint** item in the list; if you want to create a different **Viewpoint**, you can select another one directly without necessarily selecting the **No Viewpoint** item first.

You also have the option to disable Operations, which breaks apart the request and response messages so that they appear in their original chronological order, similar to the way Network Monitor displays messages. You can do this by selecting the **Disable Operations Viewpoint**. The result has similarities with the data view that is achieved when you click the **Flat Message List** button, which also simulates the Network Monitor display as described in Creating a Flat Message List.

More Information

To learn more about **Viewpoints**, see the Applying and Managing Viewpoints topic.

Viewing Message Details

You can obtain a full visual representation of message details in the **Analysis Grid** viewer, including field names, values, and types, by double-clicking any top-level parent message node or nested child message node. The indicated information is presented inline on a **Fields** tab. Note that the inline data can also include other data tabs such as the **Stack**, **Diagnosis**, and **Embedded** tabs, which provide other related message information that is described in the Message Details Tool Window topic. The figure below shows message field details inline on the **Fields** tab that displays when you double-click a message in the **Analysis Grid** viewer.

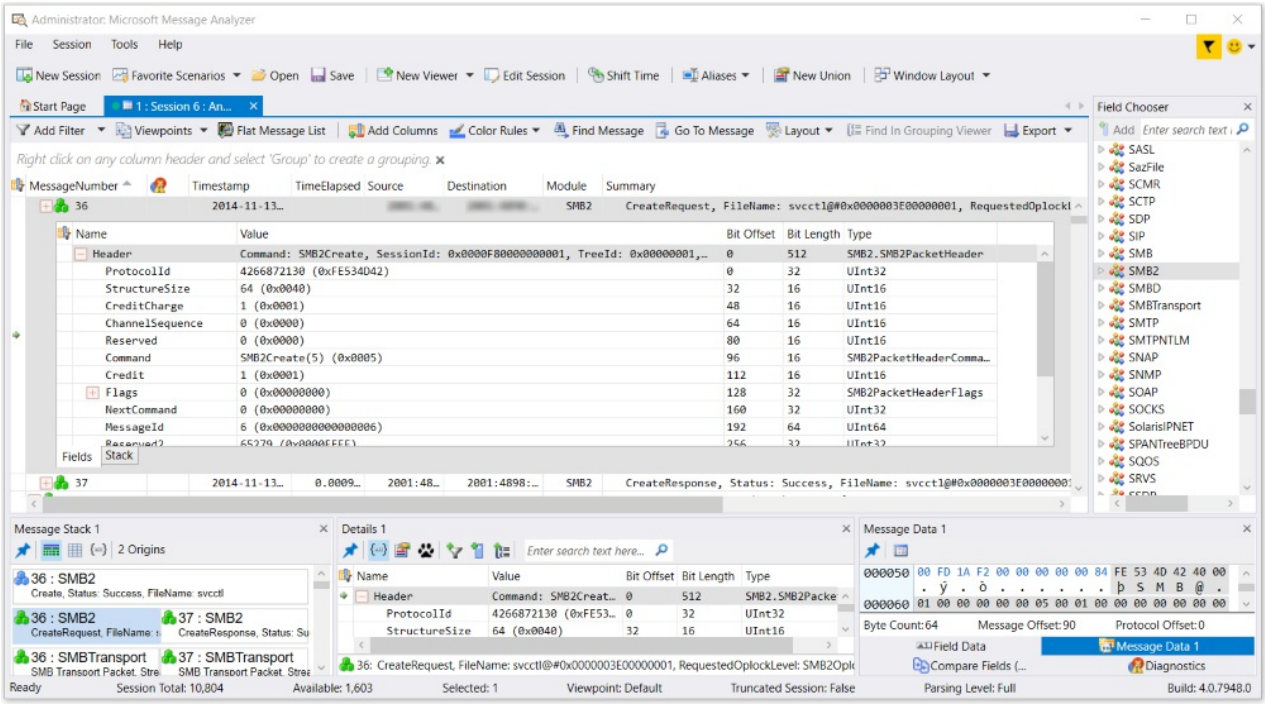

Figure 9: Message Analyzer inline message Details

You can also select any message in the **Analysis Grid** viewer to see the identical field details data in a separate window that is called the **Details** window, which typically displays below the **Analysis Grid** viewer and includes field **Name**, **Value**, **Bit Offset**, **Bit Length**, and **Type** data. For example, by selecting an **Analysis Grid** viewer message, the **Details** window immediately snaps to the selection and presents the field data for the selected message. Note that any field that you select in the **Details** window can drive the display of a hexadecimal value in the **Message Data** window or a decimal value in the **Field Data** window.

Viewing Other Message Data

Other **Tool Windows** are also available to enhance your data analysis perspective, for example, the **Message Data**, **Field Data**, **Diagnostics**, and **Decryption Tool Windows**. You can also view stack information in a separate window known as the **Message Stack Tool Window**, which provides an alternate view of the origins tree (message stack) below any top-level message that is normally hidden by collapsed message nodes in the **Analysis Grid** viewer. Note that many Message Analyzer **Tool Windows** are interactive, because they either drive or are driven by message or data selection in other windows or data viewers. For instance, by selecting a field in the **Details Tool Window**, the **Message Data** window immediately snaps to the selection and highlights the corresponding hexadecimal value of the selected field.

More Information

To learn more about Message Analyzer **Tool Windows**, see the Tool Windows topic.

Viewing Data from Multiple Sessions

Message Analyzer also provides session viewer navigation functionality from the **Session Explorer Tool Window**, to enable you to easily explore the data in different types of session data viewers, which can include **Chart** viewer **Layouts** that employ top-level data summaries in various graphic and tabular formats, the **Grouping** viewer, a **Pattern Match** viewer, the **Gantt** viewer, and several others that Message Analyzer provides. The **Session Explorer** window is accessible from the **Windows** drop-down list in the global Message Analyzer **Tools** menu.

NOTE

By right-clicking a session node in **Session Explorer**, you are presented with the **New Viewer** context menu item, which displays a drop-down list that enables you to select other data viewers that display data in separate session viewer tabs. Thereafter, any new data viewer that you specified is listed and uniquely identified by a color code in the **Session Explorer** window navigation area. If you select any **Session Explorer** node, Message Analyzer responds by immediately displaying the data on the session viewer tab that corresponds with the selected node.

The figure that follows shows an example of the **SMB Top Commands Layout** for the **Chart** viewer that you can select from the **Charts** drop-down list in the **New Viewer** drop-down list, which is accessible from the **Session Explorer** context menu. This **Layout** enables you to obtain a high-level summary view that depicts the relative distribution of traffic volume, from the highest to the lowest volume, for SMB commands in a set of trace results. The **Layout** uses a Bar element visualizer component to display the command volumes for various operations such as **Create**, **Read**, **Close**, and **QueryInfo**, which enables you to quickly evaluate the SMB commands that are consuming the most bandwidth.

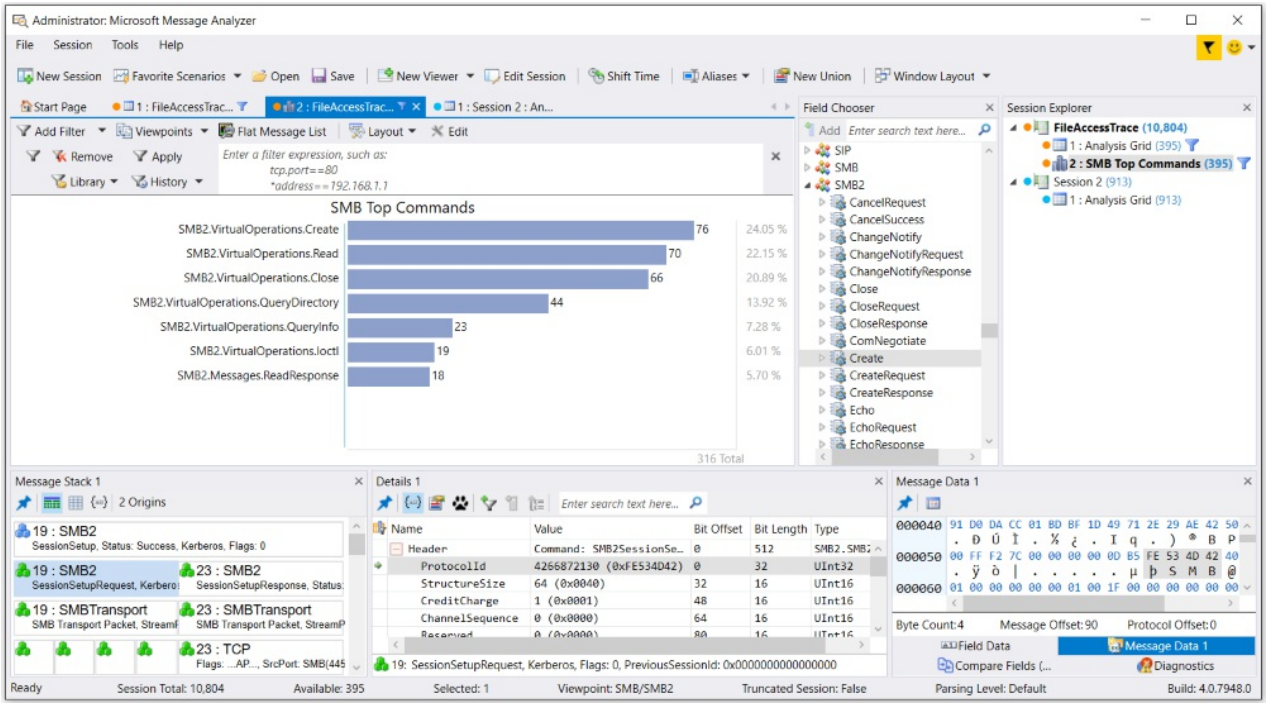

Figure 10: Message Analyzer SMB Top Commands Chart viewer Layout

In the figure, note that the **Session Explorer** window uses a common color code to identify session and viewer nodes of the same session. The same color code is used to correlate the corresponding session viewer tabs above the main analysis surface, for ease of identification. Also, session viewer nodes in **Session Explorer** and session viewer tabs are assigned a *different* color code to distinguish the data of different sessions. In addition, if any assets have been applied to a session, such as a view **Filter** or **Viewpoint**, a funnel icon displays to the right of the session viewer node in **Session Explorer**. Also note that when hovering your mouse over a session node in **Session Explorer** or over a session viewer tab, a tool tip appears with additional information, for example, the available message count and/or type of asset applied to the session viewer.

Limiting the Scope of Applied Assets

The effects of assets that you apply to any data viewer are limited in scope to the data viewer where you apply the asset. This means that no other data viewer will be affected by this action, whether the viewer is in the same session or a different session. You should note that the Filtering toolbar, from where you apply assets such as view **Filters**, **Time Filters**, and **Viewpoints**, is displayed above every data viewer that contains trace results. This is the case for **Chart** viewer **Layouts** as well. This enables you to apply different assets to different data viewers without the effects extending outside a particular viewer where an asset is applied. Note that the **Grouping** viewer has a

separate instance of the same Filtering toolbar and any assets that you apply to the **Grouping** viewer affects the **Grouping** viewer display only.

Driving Interaction Between Data Viewers

Some Message Analyzer data viewers are *interactive*, in that data selection in one viewer drives the display of data in another viewer (or **Tool Window**). For example, in a **Chart** viewer **Layout**, you can double-click a bar element in the **Bar** visualizer component or a module node in the **Timeline** visualizer component that represents the messages of a particular protocol that were captured in a trace, and display only those messages in a new **Analysis Grid** viewer tab for data assessment purposes. You might do this, to isolate a group of messages where further investigation is required. Similarly, you can select a message in the **Analysis Grid** viewer and drive the display of the network stack in the **Message Stack** window.

Other types of interactions that occur when performing actions such as message, field, or session selection, include the following.

- Selecting individual messages in a viewer such as the **Analysis Grid** and displaying corresponding message details in the **Details** window.
- Selecting a message Group in the **Grouping** viewer and displaying/isolating those messages in the **Analysis Grid** viewer. Altering the data display in a **Chart** viewer **Layout** by message Group selection.
- Selecting a message row in the **Diagnostics** window, which highlights the corresponding message/s in the **Analysis Grid** viewer.
- Selecting a field in the **Details** window that drives the highlighting of hexadecimal values in the **Message Data** window.
- Selecting a session, which refreshes the data in the **Diagnostics** window, **Message Data** window, and the **Decryption** window (if displayed with decryption results data).

Using Window Layouts

Message Analyzer enables you to customize the working environment in which you manipulate data and perform analysis. Message Analyzer does this by providing several built-in **Window Layouts** that organize the **Analysis Grid** viewer along with different **Tool Windows** into preset configurations that enable you to customize your working environment for the type of troubleshooting and analysis you perform. The window layouts are accessible from the **Window Layout** drop-down list on the global Message Analyzer toolbar. When you shut down Message Analyzer, the window configuration that you last displayed is registered in a configuration file so that the window configuration persists through subsequent Message Analyzer startups.

The **Window Layout** presets that you can select range from simple to increasingly more complex selections, given that they are intended to accommodate a cross-section of typical Message Analyzer users. The typical layout configuration consists of a single/default data viewer and an arrangement of one or more **Tool Windows**. However, you can organize your data windows any way you want.

By default, Message Analyzer uses the **Analysis Grid** viewer in all the built-in **Window Layouts**; however, after you display one of the presets, you can select a different viewer of choice if you wish. You can also add other **Tool Windows** to any of the built-in **Windows Layouts**, as needed, although you cannot modify the configuration of the built-in **Window Layouts**. Rather, any **Tool Windows** that you add to a displayed **Window Layout** are registered in the previously mentioned configuration file to persist the configuration across Message Analyzer restarts.

More Information
To learn more about **Window Layouts**, see Working with Message Analyzer Window Layouts

---

Using Message Analyzer Profiles

Message Analyzer now provides the **Profiles** feature, which enables you to use built-in or custom-specified data viewer and **Layout** presets that activate whenever you load data from specific types of input files. Prior to the

introduction of this feature, you had to manually select viewer **Layouts** in which to display your data, whenever you wanted to analyze data from different types of input files that you load into Message Analyzer. Although, it is likely that you had to engage in a trial-and-error process to discover the best **Layout** with the right context for the type of data you are analyzing. Even then, earlier versions of Message Analyzer had a minimal selection of **Layouts** from which to choose, but this is remediated in Message Analyzer v1.4.

Because Message Analyzer viewing components can expose data in different ways, you can obtain different analysis contexts for the data with different viewer **Layouts**, although if you are a new user, you may not always know which viewer **Layout** will maximize your data analysis capabilities in a given instance. The default **Layout** for the **Analysis Grid** viewer contains a baseline set of data columns that is suitable for many environments, as described in the Default View Layout topic. However, this is only a starting point, as there are many different **Layouts** that you can select from the **Layout** drop-down list on the **Analysis Grid** viewer toolbar. Similarly, you can select numerous **Layouts** for the **Grouping** and **Chart** viewers.

Displaying Predefined Analysis Environments with Built-in Profiles Some of the **Layouts** that Message Analyzer provides for the previously indicated data viewers are designed to work with each other to create an integrated and interactive analysis environment that exposes key information. You can select these manually if you know which ones are designed for integrated analysis, or to automate the process, you can simply select one of many built-in Message Analyzer **Profiles** that each define different **Layout** configurations for the **Analysis Grid**, **Chart**, and **Grouping** viewers, depending on the type of input data to be analyzed.

After a specific **Profile** is enabled in the **Options** dialog, as shown in the next figure, its preset viewer and **Layout** configuration automatically displays with populated data whenever you load data from an input file type for which the enabled **Profile** is designed, for example, a *.etl, *.cap, or *.log file. The analysis environments created by the built-in **Profiles** are predefined by Microsoft to expose the data that is typically the most important for problem solving and to expose it in a way that provides multiple perspectives on the data, from low-level details and calculated statistics to high level overviews and other data summaries.

The built-in **Profiles** along with usage overviews and analysis examples are described in Working With Message Analyzer Profiles.

> **NOTE**
>
> Message Analyzer **Profiles** are contained in an updatable package that is known as the **Message Analyzer Profiles** asset collection. You can set this asset collection for automatic updates in the **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu.

Locating the Built-In Profiles

The figure that follows shows the **Profiles** tab of the **Options** dialog, where the **Advanced Profiles** list contains all the built-in **Profiles** that are available for selection/enabling. You can use these **Profiles** as is, or you can create your own **Profiles** with the use of the **Add Profile** feature. If you want to see the internal configuration of viewer **Layouts** for any of the built-in **Profiles**, select the **Profile** of interest and then click the **Edit Profile** button on the **Advanced Profiles** toolbar. Note that the built-in **Profiles** are ReadOnly and cannot be edited, although you can edit any **Profile** that you custom design. You can access the **Options** dialog from the global Message Analyzer **Tools** menu.

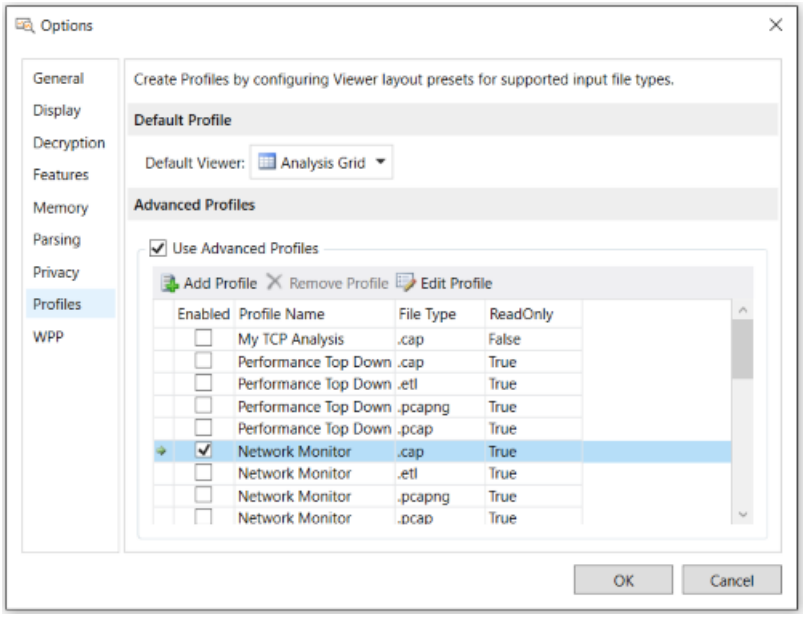

Figure 11: Message Analyzer Profiles tab of the Options dialog

Example Scenario

A scenario in which you could use a built-in **Profile** might be if you regularly analyze *.cap files for specific types of information that require a particular view of data that quickly exposes the information you need to examine for capture file analysis. To display a typical viewer and **Layout** configuration for data in this file type, Message Analyzer enables you to use the built-in **Network Monitor Profile** for *.cap files, which defines a data viewer and **Layout** configuration that is suitable for analysis of capture file data. When this **Profile** is *enabled* and you load data from a *.cap file, Message Analyzer will automatically populate the data in the viewer and **Layout** configuration that is described in the table that follows. You can view this configuration of viewers and **Layouts** in the **Network Monitor Profile** dialog that displays when you click **Edit Profile** while the **Network Monitor Profile** for *.cap files is selected in the **Advanced Profile** list on the **Profiles** tab of the **Options** dialog.

Table 2. Viewer/Layout Configuration for the Network Monitor Profile

VIEWER	LAYOUT	TYPE
Analysis Grid	Network Monitor	Tree grid
Grouping	Process Name and Conversations	Nested group
Chart	Top TCP/UDP Conversations by Message Count	Table grid

The figure that follows shows what this viewing configuration looks like after data from a .cap file is loaded into Message Analyzer.

> **NOTE**
>
> You will need to manually open the **Chart** viewer **Layout** for the **Profile** by selecting the **Default** item in the **Chart** drop-down list in the **New Viewer** drop-down list on the global Message Analyzer toolbar. Whenever the data of an input file related to an enabled **Profile** is loaded into Message Analyzer, selecting the **Default** item references the **Chart** viewer **Layout** configured in the **Profile** and causes it to be displayed.

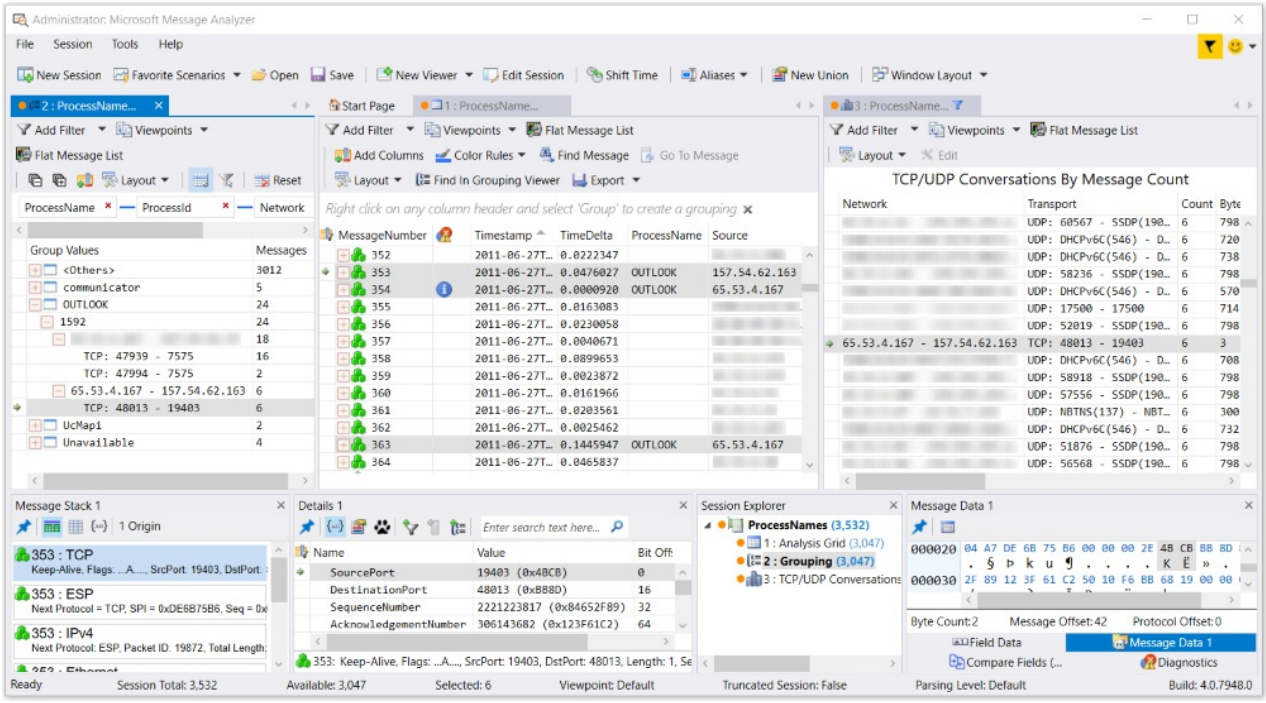

Figure 12: Message Analyzer Network Monitor Profile components

In the figure, an IP conversation is selected under the Outlook **ProcessName** group in the **Grouping** viewer, which is on the left side of the Message Analyzer user interface (UI). Because the **Grouping** viewer is in **Selection Mode**, as described in Grouping Viewer Modes of Operation, Group selection causes the messages that correspond to the conversation to be interactively highlighted in the **Analysis Grid** viewer — of which the preceding figure shows only three due to display constraints. By identifying these messages, you can then analyze them in further detail with the use of the **Message Stack**, **Details**, and **Message Data Tool Windows**.

In addition, the same conversation is selected in the **TCP/UDP Conversations by Message Count** view **Layout**, which uses a **Table** grid visualizer component to provide a data set that includes statistics such as conversation message count, payload, data transmission rate, and duration. Note that many of the data column values, such as **Count**, **Bytes**, **KBs**, **Duration**, and **BPS**, are calculated values based on data formulas that were created by Microsoft with the **Edit Chart Layout** dialog.

More Information

To learn more about using the **TCP/UDP Conversations by Message Count** Chart viewer **Layout**, see the TCP/UDP Conversations by Message Count topic.

In Conclusion

The built-in Message Analyzer **Profiles** are important tools for data correlation, analysis, and problem solving. They enable you to display integrated analysis environments that expose key data fields, calculated statistics or other low-level details, and data summaries that help you to achieve the data perspectives you need to quickly discover areas where issues are occurring. If you configure your own custom-designed **Profile/s** you have the opportunity to decide which viewers and **Layouts** you will use to expose your data.

More Information

To learn more about Message Analyzer **Profiles**, see Working With Message Analyzer Profiles.

To learn more about the Message Analyzer **Session Explorer Tool Window**, see the Session Explorer Tool Window topic.

To learn more about using the **Asset Manager** dialog, see the Asset Manager topic.

To learn more about the Message Analyzer data viewer infrastructure, see Data Viewer Concepts.

To learn more about Message Analyzer data viewers that you can work with during data analysis, including numerous **Layouts** for the **Chart** viewer, see the Data Viewers topic.

Filter Message Data

Message Analyzer provides numerous filtering capabilities to enhance data retrieval, capture, and assessment processes. Filtering is critical for focusing on specific messages and enhancing performance. For example, if you were unable to filter message data in a Live Trace Session, you might need to examine potentially tens of thousands of messages to isolate a specific problem. What most Message Analyzer users need to observe is usually related to a specific protocol, error message, conversation, or process. By providing the ability to filter while retrieving, capturing, or viewing data, Message Analyzer provides a convenient way to reduce the scope of the data that you are working with and more effectively pinpoint your issues.

The material that describes these capabilities is included in the sections that follow.

Using a Session Filter
Using Special Filters for a Live Trace
Using View Filters to Manipulate a Set of Trace Results
Creating Custom Filters

Using a Session Filter

When capturing data or loading data into Message Analyzer through a Live Trace Session or a Data Retrieval Session, as shown in the figures of the earlier sections: Configuring a Live Trace Session and Retrieve Message Data, you can use the **Session Filter** feature to isolate specific data that you want to work with. You can select a built-in **Session Filter** from the **Message Analyzer Filters** asset collection **Library** drop-down list that appears on the **Session Filter** toolbar of the **New Session** dialog, or you can create a custom **Filter** of your own design. After specifying a **Session Filter** and clicking the **Start** button for a configured Live Trace Session or Data Retrieval Session, the filtering action is automatically applied in the background as messages are filtered and delivered to the default data viewer, for example, the **Analysis Grid** viewer. A **Session Filter** works in the same way most filters work, by passing data that matches the filtering criteria and dropping any data that does not. However, you should carefully note that you can never recapture the data that you filter out with a **Session Filter** in a Live Trace Session, whereas with a Data Retrieval Session, you can always click the **Edit Session** button on the global Message Analyzer toolbar to return to session configuration, where you can remove or recast your filtering criteria and then reload the data from the originally specified saved files. A **Session Filter** is shown in the figure that follows.

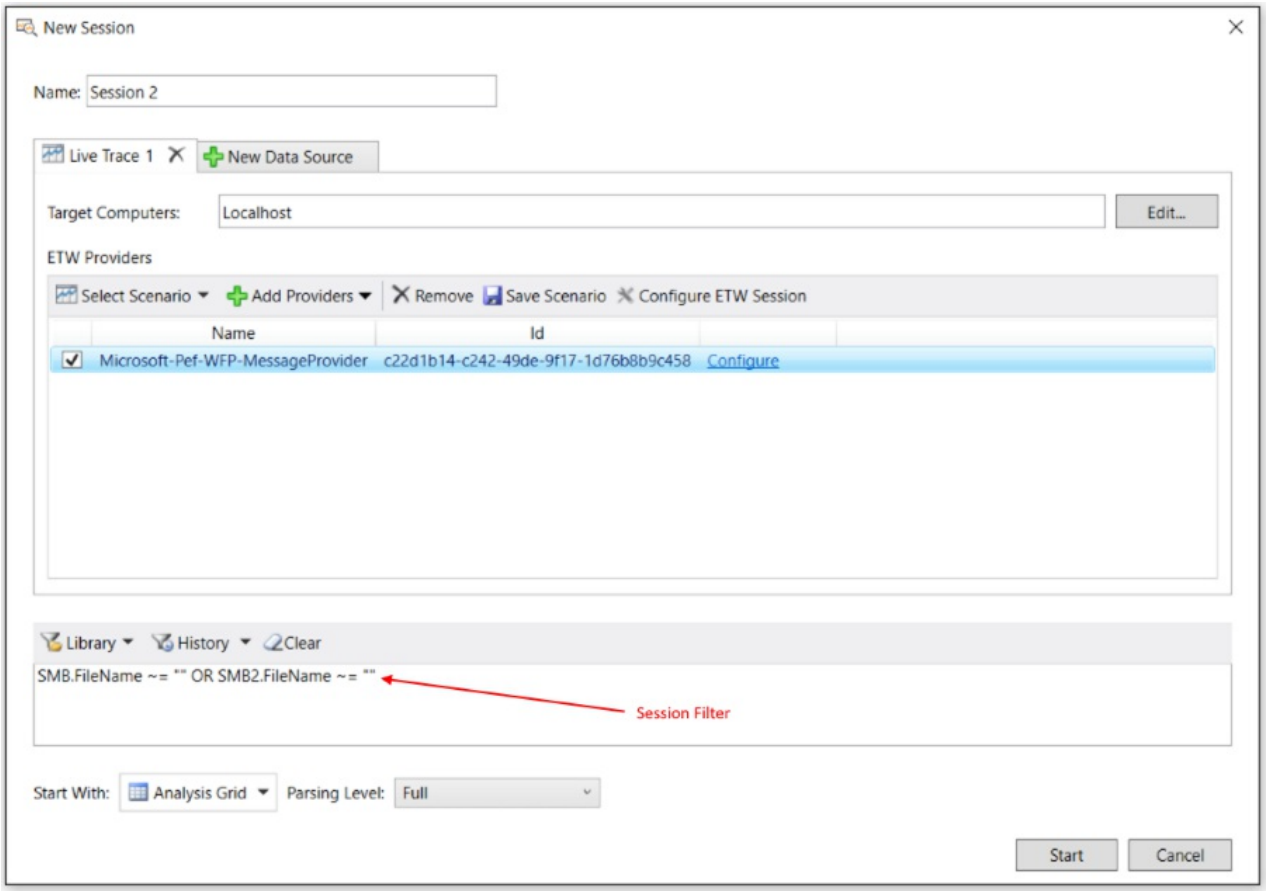

Figure 13: Session Filter for a Live Trace Session

For instance, when configuring a **Session Filter**, you could specify a Filter Expression that isolates messages to a specific network address, port, or protocol, or that contains a particular field value or other text. For a Live Trace Session, the effects of a **Session Filter** are applied at the time of data capture, therefore, your trace results will already reflect application of the filtering. For a Data Retrieval Session, the effects of a **Session Filter** are applied at the time of data loading, therefore, the loaded data will already reflect application of the filtering. By contrast, the effects of a view **Filter** are applied to a set of trace results or loaded data *results* and are temporary, as you can alternately remove or apply the **Filter** repeatedly as required, or even modify it, during data analysis.

The figure that follows reflects the application of the above specified **Session Filter** `SMB.FileName ~= "" OR SMB2.FileName ~= ""` during data capture, which limits the trace to SMB or SMB2 messages that have a **FileName** field populated with data, as described in the File Sharing Category topic. Thereafter, an **IPv4 Gradient Right Color Rule** (with dark green highlights) was applied to the Live Trace Session results to quickly expose messages that are using the IPv4 protocol, for analysis purposes.

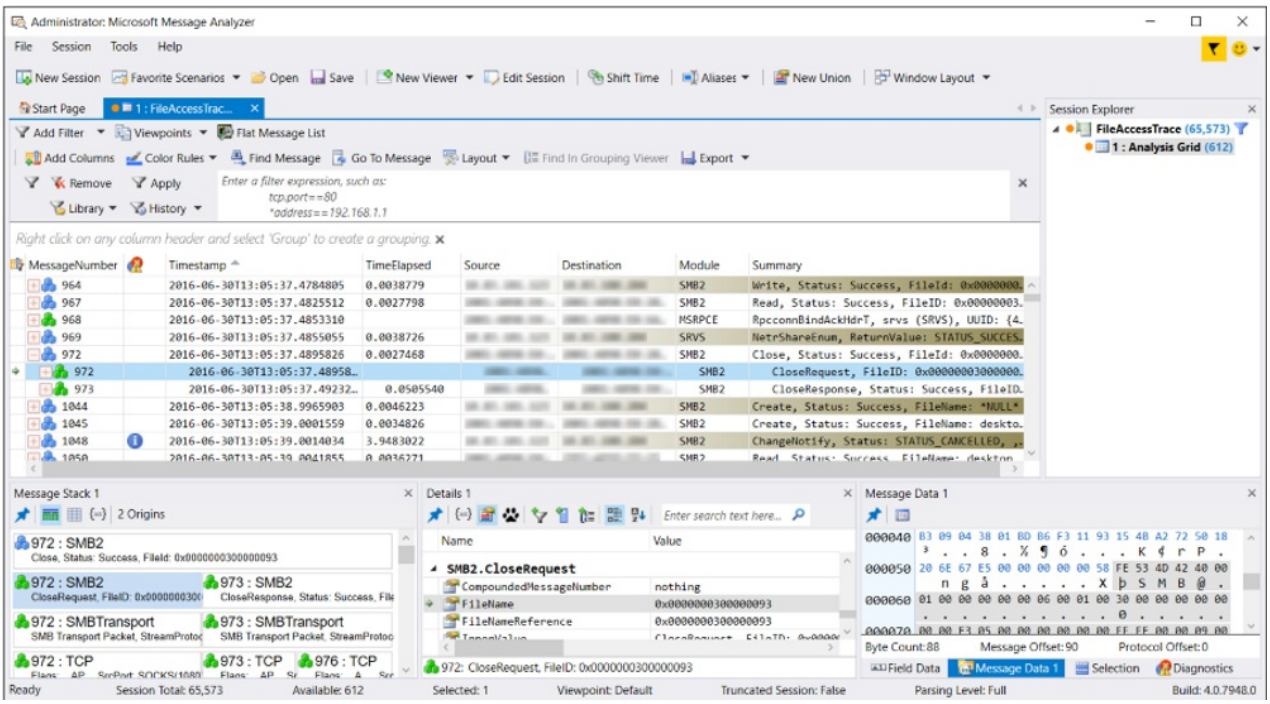

Figure 14: Message Analyzer Analysis Grid viewer results with a Session Filter and Color Rule Applied

Using Special Filters for a Live Trace

You also have the option to use many other types of filters in a Live Trace Session, depending on the **Trace Scenario** and operating system you are running, as follows:

- **Fast Filters** and **WFP Layer Set** filters — accessible from the **Provider** tab of the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** configuration dialog. You can display this dialog by clicking the **Configure** link to the right of the **Microsoft-PEF-WFP-MessageProvider** listing on the **Live Trace** tab of the **New Session** dialog, after you select one of several **Trace Scenarios** that contain this provider from the **Select Scenario** drop-down list on the **ETW Providers** toolbar of the **Live Trace** tab. For example, the **Network Tunnel Traffic and Unencrypted IPSEC**, **Loopback and Unencrypted IPSEC**, and **Local Loopback Network Trace Scenarios** all use the **Microsoft-PEF-WFP-MessageProvider**.

- **Fast Filter Groups** and **Adapter filters** — accessible from the **Provider** tab of the **Advanced Settings – Microsoft-PEF-NDIS-PacketCapture** configuration dialog. You can display this dialog by clicking the **Configure** link to the right of the **Microsoft-PEF-NDIS-PacketCapture** listing on the **Live Trace** tab of the **New Session** dialog, after you select the **Local Network Interfaces Trace Scenario** from the **Select Scenario** drop-down list on the **ETW Providers** toolbar of the **Live Trace** tab. The **Microsoft-PEF-NDIS-PacketCapture** provider is available on computers running the Windows 7, Windows 8, or Windows Server 2012 operating system only.

- **HostName** and **Port** filters — accessible from the **Provider** tab of the **Advanced Settings – Microsoft-Pef-WebProxy** configuration dialog. You can display this dialog by clicking the **Configure** link to the right of the **Microsoft-Pef-WebProxy** provider listing on the **Live Trace** tab of the **New Session** dialog, after you select the **Pre-Encryption for HTTPS Trace Scenario** from the **Select Scenario** drop-down on the **ETW Providers** toolbar of the **Live Trace** tab.

- **Event Keyword** and error **Level** filters — accessible from the **ETW Core** tab of any **Advanced Settings** dialog for any **Trace Scenario** that you select from the **Select Scenario** drop-down list. You can display this dialog by clicking the **Configure** link to the right of any provider listing on the **Live Trace** tab of the **New Session** dialog. Note that not all ETW Providers contain an event **Keyword** configuration.

- **NDIS stack, Hyper-V-Switch extension layer, and Host adapter filters** — accessible from the **Provider** tab of the **Advanced Settings – Microsoft-Windows-NDIS-PacketCapture** provider

configuration dialog. You can display this dialog by clicking the **Configure** link to the right of the **Microsoft-Windows-NDIS-PacketCapture** provider listing on the **Live Trace** tab of the **New Session** dialog, after you select the **Local Network Interfaces**, **Remote Network Interfaces**, or **Remote Network Interfaces with Drop Information Trace Scenario** from the **Select Scenario** drop-down list on the **ETW Providers** toolbar of the **Live Trace** tab. The **Microsoft-Windows-NDIS-PacketCapture** provider has remote capabilities and is available on computers that are running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system only.

The filters that are available for the **Microsoft-Windows-NDIS-PacketCapture** provider in these scenarios consist of advanced driver-level filters that include the following:

- Host adapter filters

> **NOTE**
> If you want to isolate traffic to a particular virtual machine (VM) that is serviced by a Hyper-V-Switch, you should select the VM adapter in the Interface Selection (upper) section of the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog to select the adapter and then specify the MAC address of the VM adapter in the **MAC Addresses** box in the **Filters** section of the dialog, rather than simply selecting (enabling) the adapter. Otherwise, you could return all switch traffic rather than the traffic of a selected VM, given that a Hyper-V-Switch driver cannot distinguish between VMs.

- NDIS stack and Hyper-V-Switch extension layer filters
- Truncation filters
- Packet traversal direction filters
- IP protocol number filters
- MAC address filtering
- IP address filters

More Information
To learn more about the filtering capabilities of the **Microsoft-Windows-NDIS-PacketCapture** provider, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.
To learn more about **Fast Filters** and **WFP Layer Set** filters, see the Microsoft-PEF-WFP-MessageProvider topic.
To learn more about **Fast Filter Groups** and **System Network Adapter** Group filters, see the PEF-NDIS Fast Filters and Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog topics.
To learn more about **HostName** and **Port** filters, see the Microsoft-PEF-WebProxy Provider topic.
To learn more about **Keyword** event and error **Level** filters, see the System ETW Provider Event Keyword/Level Settings topic.
To learn more about NDIS stack, Hyper-V-Switch extension layer, host adapter, and other special filters, see the Configuring a Remote Capture and Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog topics.

- **Time Filter** — you can utilize a **Time Filter** to configure a window of time in which to view the results of a Message Analyzer session. This is particularly useful if you can approximate a time frame in which you suspect a particular issue occurred that you need to detect. It is also useful in situations where you are working with data from multiple input files in different time zones to which you have applied a **Time Shift** and you want to view all the data that exists in a particular time slot, for which you can create a **Time Filter** window.

  Applying a Time Filter to Captured Data The major advantage of using a **Time Filter** against a set of

trace results is that you can remove it, modify the time window, reapply it, and repeat this process as many times as needed. You can display the **Time Filter** configuration controls for session results by selecting **Add Time Filter** from the **Add Filter** drop-down list on the Filtering toolbar that appears above all session data viewers. The **Time Filter** configuration panel contains the time slider controls and window definition readouts, along with the **Apply** and **Remove** buttons that enable you to alternately apply a **Time Filter** or remove its effects as required.

Applying a Time Filter to Loaded Data You can also utilize a **Time Filter** to configure a window of time in which to view static data that you load into Message Analyzer from selected input files. Note in this case that the **Time Filter** configuration is applied *as* you load the data, rather than *after* you load the data. To apply a **Time Filter** to data that you are loading, you must create the **Time Filter** window during Data Retrieval Session configuration in the **New Session** dialog, where the **Time Filter** controls are located below the input **Files** list in the dialog. The **Time Filter** configuration will be applied by Message Analyzer after you click the **Start** button in the **New Session** dialog.

Note that you have the option to modify the original **Time Filter** configuration in the **Edit Session** dialog for the Data Retrieval Session by clicking **Edit Session** on the global Message Analyzer toolbar. You will need to set the **Edit Session** dialog for **Full Edit** mode to enable the **Time Filter** controls. When the adjustments are complete, click the **Apply** button in the **New Session** dialog to apply the **Time Filter** changes to the data set.

Also note that if you did not apply a **Time Filter** to the data loading process in a Data Retrieval Session, you still have the option to utilize the **Time Filter** feature from the **Add Filter** drop-down list on the Filtering toolbar. From this location, you can use simple button clicks to alternately **Apply** and **Remove** the filtering effects as required.

IMPORTANT

Message Analyzer provides you with the versatility to apply a **Time Filter** to the results of a Live Trace Session, the results of a Data Retrieval Session, or to the data loading process. In the latter case, you can achieve performance enhancements due to the effects of a **Time Filter** on reducing the input message volume that is loaded into Message Analyzer. But this can have an effect on usability when the filtered-out messages have a bearing on the analysis in which you are engaged. When this is the case and you want to recover messages that the input **Time Filter** dropped, you will need to edit the session as described earlier, to create a different **Time Filter** configuration; this also has an impact on usability. Therefore, you might want to further consider the tradeoffs between performance and usability, especially when loading data from very large files.

More Information

To learn more about the impacts on performance and usability with the **Time Filter** feature, see Considering Performance vs. Usability Factors for Time Filter Application.

More Information

To learn more about **Time Filters**, see Applying a Time Filter to Session Results and Applying an Input Time Filter to a Data Retrieval Session.

Using View Filters to Manipulate a Set of Trace Results

After you capture or retrieve your message data in a Live Trace Session or Data Retrieval Session, respectively, you have a baseline set of trace results to work with. However, it is very likely that to analyze the data, you will want to manipulate it with various Message Analyzer tools to isolate specific messages of interest that can expose issues you are trying to detect. One of the most common ways to do this, is to use a view **Filter** to filter for data that is relevant to the problem you are trying to solve while filtering out data that isn't. This enables you to create a set of messages that is focused on the data you need to examine, without the encumbrance of scrutinizing potentially hundreds if not thousands of messages that are irrelevant to the issue at hand. When you apply a view

Filter, the original data set is always preserved and re-displays after you remove it. Note that the effects of a view **Filter** apply to the in-focus data viewer only and do not impact other viewers, even in the same session.

You can display the configuration controls for a view **Filter** by selecting **Add Filter** from the **Add Filter** drop-down list on the Filtering toolbar that appears above all session data viewers. The controls that display in the **Filter** configuration panel enable you to specify a built-in or custom **Filter**, and then apply and remove it as required, as described in Applying and Managing Filters.

The built-in view **Filters** are contained in a centralized **Library** that is exposed in the following locations.

- **Filter** configuration panel that appears when you select **Add Filter** in the **Add Filter** drop-down list on the Filtering toolbar above any session data viewer.
- **Filter** configuration panel that appears when you select **Add Filter** in the **Add Filter** drop-down list on the Filtering toolbar above the **Grouping** viewer toolbar.
- **Session Filter** toolbar in the **New Session** dialog; for use when you are configuring a new Live Trace Session or Data Retrieval Session.
- **Add Viewpoint Filter** panel that appears when you select the **Add Viewpoint Filter** item from the **Add Filter** drop-down list on the Filtering toolbar. Note that a **Viewpoint** must be already applied to a set of trace results for this list item to be available for selection. Helps you to further refine your analytical focus on specific messages.
- **Find Message** panel that is accessible from the **Analysis Grid** viewer toolbar.
- **Edit Color Rule** dialog, which is accessible by clicking the **New Color Rule** in the **Color Rule** drop-down list on the **Analysis Grid** toolbar. Note that typical configuration of a **Color Rule** includes specifying a Filter Expression from the centralized **Library**.

> **TIP**
> You can also specify a view **Filter** for a set of trace results by right-clicking a data field value in an **Analysis Grid** viewer column and selecting **Add '<ColumnName>' to Filter**, where "<ColumnName>" is a placeholder for the data column under which the data field appears. Note that this action automatically creates the Filter Expression in the **Filter** configuration panel, but does not apply it. As a result, you must manually apply such a **Filter** by clicking the **Apply** button in the **Filter** configuration panel. This feature enables you to automatically code a column value into a valid Filter Expression, which you can quickly apply to a set of trace results.

To specify a view **Filter**, **Session Filter**, **Find Message** filter, **Color Rule** filter, or **Viewpoint Filter** for a set of trace results, you will need to either select a built-in Filter Expression from the centralized **Library** in the above specified locations, or manually create one as described in Writing Filter Expressions. You will then need to click the **Apply** button (or **Find** command in the case of **Find Message** filters) for the **Filter** configuration to take effect. The centralized **Library** contains the built-in Filter Expressions that are provided by the **Message Analyzer Filters** asset collection in every Message Analyzer installation, for which you can use the following for the indicated purpose:

- **Asset Manager** dialog — to manage downloads and auto-sync updates for the **Message Analyzer Filters** asset collection or other collections. Asset Manager is accessible from the global Message Analyzer **Tools** menu.
- **Manage Filters** dialog — to export and import asset collection items to and from others, respectively, for mutual sharing. The **Manage Filters** dialog is accessible from every user **Library** drop-down list.

More Information
To learn more about the functionality of the built-in view **Filters**, see the Filtering Live Trace Session Results topic, which describes each **Filter** in the centralized Filter Expression **Library**.

To learn more about auto-syncing, downloading, and managing the **Message Analyzer Filters** asset collection with the **Asset Manager** dialog, see the Sharing Infrastructure and Managing Asset Collection Downloads and Updates topics.

Creating Custom Filters

To create your own Filter Expressions, you will need to understand the Message Analyzer Filtering Language. This Operating Guide devotes a significant amount of coverage to the subject, to help you understand and use the Filtering Language, as described in the "More Information" section that follows. Note that Message Analyzer provides the Filter IntelliSense service to assist you in creating your own Filter Expressions. Filter IntelliSense is an interactive and intelligent statement completion service that responds to the text that you enter in any Filter Expression text box, by providing a display of choices in response to the characters you type.

When you create your own custom **Filters** you must save them to the centralized Filter Expression **Library** that is exposed in the locations described earlier, that is, if you want such **Filters** for future use and for sharing with others. However, before you save a **Filter** that you created, Message Analyzer performs a simple verification check to ensure that you have a valid expression, although checks on field names are less restrictive in Message Analyzer v1.4 to enable operation with other parsers. Note that when you create and save a custom **Filter**, it is located to the **My Items** category in the Filter Expression **Library**. Thereafter, you can simply select your custom **Filter** from the **Library** whenever you want to use it.

More Information
To learn more about the Filtering Language, see Writing Filter Expressions.
To learn more about the Filter IntelliSense feature, see Filter IntelliSense Service.

Analyze Message Data

When analyzing data that you have either captured live on the network, loaded into Message Analyzer, or retrieved from a device such as a Bluetooth, you have the option to apply various types of filters to manipulate the way data is presented for analysis purposes. For example, you could apply various view **Filter**, **Time Filter**, **Color Rule**, **Column Filter**, and **Grouping** configurations to a set of trace results, to name a few. In addition, you might use the **Pattern Match** capability to detect message patterns across a set of trace results.

Advisory
To review summary descriptions of the analysis tools that are available in Message Analyzer, see Analyzing Message Data. For further details about the tools mentioned in this topic, see "More Information" at the end of this section.

Data Analysis Feature Highlights
Some highlights of the options you have for manipulating data are included here in the following features.

- **Viewpoints** — you have the option to apply a **Viewpoint** to enhance your data analysis and troubleshooting perspectives. When a **Viewpoint** is applied, you can examine network traffic from the perspective of a protocol because all messages above the "viewpoint" protocol are temporarily removed from display. This feature is advantageous when you have higher-layer traffic that obscures the underlying messages that you want to troubleshoot. To apply a **Viewpoint**, select one from the **Viewpoints** drop-down list on the Filtering toolbar that appears above each Message Analyzer session viewer that you open.

  For example, you might apply a **TCP Viewpoint** to display TCP messages at top-level for diagnostic purposes. You could then select one of the **TCP** view **Layouts** for the **Analysis Grid** viewer or **Chart** viewer to expose additional data field values, calculated statistical values, or high-level data summaries that are particularly important to your analytical proceedings. You can do the same thing with an **HTTP Viewpoint** and then select one of the **HTTP** view **Layouts** for these same data viewers.

- **Grouping** — you have the option to organize trace results data into Groups that expose messages in a nested Group configuration that you can specify by executing the **Group** command that displays when you right-click an **Analysis Grid** viewer column. You can also utilize the **Grouping** viewer and select built-in **Grouping** viewer **Layouts** that organize data into unique Group configurations that are designed to create a specific analytical focus, where you can summarize and expose target data in grouping categories across a high volume of messages. The **Grouping** viewer also provides different modes of interaction with the **Analysis Grid** viewer, which includes the **Selection** and **Filtering** modes. For example, in the **Selection** mode, Group selection causes the Group messages to be *selected* in the **Analysis Grid** viewer; in the **Filtering** mode, Group selection causes the Group messages to be *filtered* in the **Analysis Grid** viewer.

- **Analysis Grid** viewer **Layouts** — you have the option to apply built-in view **Layouts** that contain an arrangement of data columns that are designed to assist you in data analysis and troubleshooting processes, for example, the **File Sharing SMB/SMB2**, **Network Conversation Tree with Process ID**, and **TCP Deep Packet Analysis with ABSOLUTE Sequence Number with Grouping** view **Layouts**. You can select a view **Layout** for the **Analysis Grid** viewer from the **Layout** drop-down list on the **Analysis Grid** viewer toolbar in an Analysis Session.

- **Flatten messages** — you can click the **Flat Message List** button on the Filtering toolbar to create a message display that resembles how messages appear in Microsoft Network Monitor. This action breaks apart request and response message pairs that Message Analyzer encapsulates in Operation nodes by default, which can have an impact on analysis. This results in reorganizing messages in their original chronological order, with the exception of message fragments which remain under the expandable top-level message nodes that are each designated by a green-cubed icon. To re-establish the default Message Analyzer display, click the **Flat Message List** button again. A disadvantage of creating a flat message list is that responses to specific requests may be difficult to locate in high volume traces.

- **Field Chooser** — for any given set of data that displays in the **Analysis Grid** viewer, you should be aware that there are many more columns of data that you can add to the grid beyond the default column layout, to expose the values of message fields that could be critical to troubleshooting processes. You simply find the relevant protocol or module of interest in the **Field Chooser Tool Window**, expand the protocol node, and navigate to the data field you want to add to the **Analysis Grid** viewer column **Layout**. After you double-click the field name, the field is added to the **Analysis Grid** viewer as a new column, provided that the **Analysis Grid** viewer is in focus when you do so. At this time, data should display in the new field for the particular protocol or module of interest, unless the field is not used or contains no data.

  If the Grouping Viewer is in focus, then double-clicking a field name in **Field Chooser** window will add a new nested Group to the current **Grouping** viewer **Layout** configuration. Note that you can also right-click any field in **Field Chooser** window and choose the **Add as Column** or **Add as Grouping** context menu command to add a column to the **Analysis Grid** viewer or a new nested Group to the **Grouping** viewer, respectively.

  > **NOTE**
  >
  > To access the **Field Chooser** window, if it is not already displayed, click the **Add Columns** button on the toolbar of the **Analysis Grid** viewer. If the **Grouping** viewer is in focus, you can access the **Field Chooser** by clicking the **Add Groupings** button on the toolbar of the **Grouping** viewer. The **Field Chooser** also appears in the **Windows** submenu of the global Message Analyzer **Tools** menu.

  > **TIP**
  >
  > View **Layouts** and column layouts are different terms that essentially describe the same feature or function in the **Analysis Grid** viewer. Also note that the **Grouping** viewer has its own set of view **Layouts** that are independent of view **Layouts** for the **Analysis Grid** viewer.

- **Pattern Matching** — provides a pattern matching capability that can identify sequential message patterns in a group of messages, for example virus signatures, processes in a faulty state that form a specific pattern, and other patterns such as request/response pairs. You can match message sequences by executing user-designed or built-in Pattern expressions that are provided with the **Pattern Match** viewer.

  Note that you can create a Pattern expression that is pre-populated with an initial configuration by selecting and right-clicking one or more related messages in the **Analysis Grid** viewer, such as HTTP or DNS, and then selecting the **Create Pattern** command that displays in the context menu that appears. You can also display the **Pattern Match** viewer from the **New Viewer** drop-down list on the global Message Analyzer toolbar to create a Pattern expression without any data population automation.

- **Filtering columns** — you can apply a **Column Filter** to any **Analysis Grid** viewer column, to filter your trace results according to search text that you specify for a column. You can also do the same for columns in the **Details Tool Window**, to filter for specific message field names or other data values in the **Details** window. Be aware that **Column Filters** search only on data that displays in top-level parent nodes; child nodes will not be included in the search unless you first expand them.

- **Aliases** — if you have any data fields that are difficult to work with, due to their cryptic or complex values, for example an IPv6 address, Message Analyzer enables you to convert their data values to a more user-friendly name for ease of recognition. To configure an **Alias**, right-click a field value that you want to convert and then select the **Create Alias for ' *<columnName>* '...** context menu item. This action causes the **Alias Editor** to display, from where you can configure a new **Alias**.

- **Unions**— if you have multiple data sources that relate to a common environment or service from which you have run traces or generated logs, it is not uncommon for the different data sources to specify different names for fields that have identical meaning and value types. If this occurs, you can create a **Union** to correlate the field values into a single, newly-named entity that you specify to reflect those values. This makes it easier to locate and analyze data in an interlaced set of messages from multiple sources.

- **Chart** viewer **Layouts** — Message Analyzer provides a wide range of **Layouts** that you can select for the **Chart** viewer. They are accessible from the **New Viewer** drop-down list on the global Message Analyzer toolbar. A description of the available **Layouts** are provided in the Chart Viewer Layouts topic. In general, they display high-level data summaries and data formula-generated statistics in several graphic formats to enhance your analysis contexts through various types of data overviews. Some of these **Layouts** are designed to work together with the **Grouping** and **Analysis Grid** viewers to create integrated and interactive analysis environments, as described in Working With Message Analyzer Profiles. Note that you can even create your own custom-designed **Chart** viewer **Layouts** for the environment in which you typically work, as described in Extending Message Analyzer Data Viewing Capabilities.

Other Data Analysis Features

Other techniques that you can use to analyze data consist of the following:

- **Tool Windows** — you can display additional **Tool Windows** to dramatically enhance the scope of analysis capabilities. These tools are accessible from the **Windows** submenu of the global Message Analyzer **Tools** menu.

- **Finding messages** — you can use the **Find Message** feature to locate individual messages. The **Find** command is designated by the **Find** binoculars icon in the **Find Message** window that displays when you click the **Find Messages** button on the toolbar of the **Analysis Grid** viewer. This command enables you to locate the next message that matches a specified Filter Expression, while still retaining visibility and context of all the messages in the original trace results. You can select either a built-in Filter Expression from the **Library** drop-down list in the **Find Message** window, or you can manually configure one, such as `contains "bing"` (or some other string), or `#MessageNumber==messagenumber`, to locate the next message that contains the value that you specify.

  By using this command, you can dramatically impact your data analysis experience. Although Message

Analyzer already provides the view **Filter** capability that works similarly, the disadvantage of a view **Filter** is that all messages surrounding the target message/s are hidden after view **Filter** application, unless they match the filter criteria. However, in many cases the context of the surrounding messages is key to the analysis. When this is the case, it might be better to employ a **Find** filter. A **Find** filter highlights the next top-level message that matches the filtering criteria, even if the match is to a message that is *within* the message stack of the highlighted top-level message, also known in this documentation as the origins tree.

> **TIP**
> You can also use the **Go To Message** feature to locate a specified message by its number across one or more data sources or sessions.

- **Sorting** — you can sort data columns in the **Analysis Grid** viewer in ascending, descending, or original capture order, to expose values or trends that can identify potential issues. For example, you can sort the **DiagnosisTypes** column of the **Analysis Grid** to bubble up all diagnostic messages for quick analysis.
- **Applying time shift values** — Message Analyzer provides a **Shift Time** dialog that enables you to apply a specific incremental time shift value to a message collection from a selected data source, when you *know* beforehand that a time shift is required. You can also specify a time shift for a particular message when you *discover* through analysis that a shift is required. Applying a time shift to a selected message then causes a recalculation of time stamps for all messages in a selected data source.

  You can use this feature to synchronize multiple traces that you load into Message Analyzer, for example to adjust for machine clock skew or time zone changes across traces. You might also want to simply match the **Timestamp** of one message loaded from a particular data source to that of another message loaded from a different data source.
- **Adding bookmarks and comments** — you can add **Bookmarks** and **Comments** for annotation purposes to coordinate data analysis with other team members.

  > **NOTE**
  > To quickly locate messages that have a **Comment**, you can add the **HasComments** field as a new **Analysis Grid** viewer column from the **General** node of the **Field Chooser**.
- **Find in Grouping Viewer** — for any message that displays in the **Analysis Grid** viewer, you can locate the **Grouping** viewer Group that contains it. You simply right-click a message in the **Analysis Grid** viewer and select the **Find in Grouping Viewer** context menu command. This results in a cross-correlation of data between viewers that can provide unique analysis contexts.
- **Parse As** — for any message type that displays in the **Analysis Grid** viewer, you can specify a different port on which that message type will be parsed. Enables you to reparse a set of trace results based on alternate ports that you specify for specific protocols, to accommodate for network traffic that used alternate ports for security purposes.
- **Message Details** — you can click the blue- or green-cubed icon to the left of any message to display the **Details** of that message inline.

More Information
To learn more about the details of working with filters and other data manipulation features for analysis, see the following topics:
Using the Filtering Toolbar
— Applying and Managing Filters
— Applying a Time Filter to Session Results

— Applying and Managing Viewpoints
— Working With Operations
— Creating a Flat Message List
Using the Find Message Feature
Using the Go To Message Feature
Filtering Column Data
Using and Managing Color Rules
Pattern Match Viewer
Using the Analysis Grid Group Feature
Grouping Viewer
Applying and Managing Analysis Grid Viewer Layouts
Using the Field Chooser
Using and Managing Message Analyzer Aliases
Configuring and Managing Message Analyzer Unions
Setting Time Shifts
Tool Windows

Save Message Data

After you have performed analysis of your message data, you have the option to save it in the Message Analyzer native .matp file format or in the .cap format, as described in Saving Message Data. Thereafter, if you want to work further with the data or share it with others, you can quickly load the data back into Message Analyzer through a Data Retrieval Session, or you can load the data by using the **Open** dialog, which is accessible from the global Message Analyzer **File** menu or from the global Message Analyzer toolbar.

The figure that follows illustrates the **Save/Export Session** dialog, in which you can choose the messages you want to save. You have the option to save all messages that you captured in a Live Trace Session or loaded from a Data Retrieval Session, filtered messages only, or you can select specific messages to save.

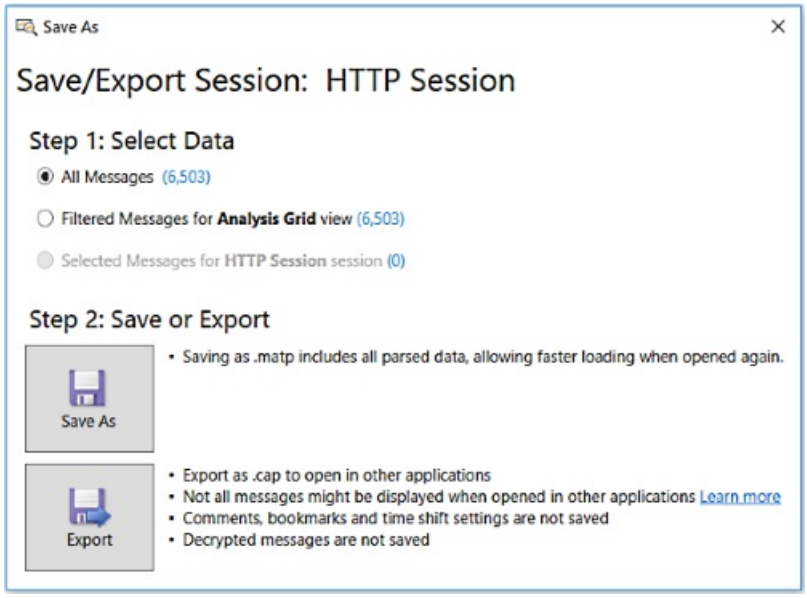

Figure 15: Message Analyzer Save/Export Session dialog

The following summarizes the different ways to save message data:

- **Use the Windows Save As dialog** — highlight one or more messages in the **Analysis Grid** viewer, right-click the group of messages, and then select the **Save Selected Messages...** context menu command to display the Windows **Save As** dialog. After you specify a name for the session messages you are saving, click the **Save** button to save the session in the native .matp file format only.

- **Use the Export feature** — click the **Export** button on the **Analysis Grid** viewer toolbar to save **All** or **Selected** messages in comma separated value (CSV) file format.
- **Use the Message Analyzer Save As dialog options** — highlight one or more messages in the **Analysis Grid** viewer and then click **Save As** in the **File** menu to display the **Save As** dialog. When you use this dialog to save data, you can specify additional save options with the use of three radio buttons under **Step 1** of the dialog, which includes the following:
  - **All Messages ()** — saves all messages in a set of trace results.
  - **Filtered Messages for <viewerName> view ()** — saves the messages that result from a filtering operation.
  - **Selected Messages for <sessionName> session ()** — saves only the messages that are selected in a set of trace results.

  To save selected messages only with the dialog, use the third option, which parenthetically indicates the number of messages that you highlighted in the in-focus data viewer. Thereafter, click the **Save As** button in **Step 2** of the dialog to open the Windows **Save As** dialog, from where you can navigate to an appropriate directory location for saving the data in the native Message Analyzer .matp file format. To export the selected messages to a .cap file, click the **Export** button in **Step 2** of the dialog to display the Windows **Save As** dialog.

If you have a session configuration that consists of an aggregation of data from multiple sources that you have analyzed, Message Analyzer enables you to save your results to a single file in the default .matp format. Note then when you export your data as a .cap file, it will be compatible with the Microsoft Network Monitor tool and other applications, with certain exceptions that are described in Compatibility with Exported CAP Files.

More Information
To learn more about saving Message Analyzer data, see Saving Message Data.

PEF Architecture Tutorial

5 minutes to read

This tutorial briefly describes the main features of the Microsoft Protocol Engineering Framework (PEF) that directly support the functions of Message Analyzer. A diagram of PEF architecture is included along with supporting conceptual descriptions, to show how Message Analyzer functions are enabled by the framework.

PEF Components

Message Analyzer is a new tool for capturing, displaying, and analyzing network traffic, system events, device messages, and log data. It is the key, outwardly-facing component in the Protocol Engineering Framework. PEF was created by Microsoft to help improve protocol design, development, documentation, and testing. The following major messaging functions are provided by various PEF components:

- Message capturing
- Message representation as Open Protocol Notation (OPN) descriptions
- Message parsing, stack and fragment reassembly, and filtering
- Message validation (data, behavior, and architecture) based on protocol-specification standards

Message Analyzer directly relies upon the following components of the PEF architecture to support its functionality:

- **OPN** — the protocol description language that enables developers to model protocol architecture, behavior, and data. The entire OPN system, including types, actors, endpoints, and flow is implemented in .NET classes. OPN and .NET classes are compiled to produce a binary representation of each OPN protocol description that defines specific protocol architecture, behavior, and data. Compiled versions of the OPN descriptions reside in a binary model cache known as the protocol object model (POM), which is in turn consulted by the Runtime component for matches to messages received on the wire.

  Message Analyzer relies upon the presence of compiled OPN protocol descriptions so it can display messages that have been captured and parsed by the PEF Runtime.

- **OPN Compiler** — provides the compilation infrastructure for OPN protocol descriptions. The OPN Compiler generates the binary structures that comprise the POM.

  Message Analyzer relies upon the OPN Compiler to ensure that all OPN definitions, descriptions, and filter expressions are verified, so that messages captured in a Live Trace Session or loaded into Message Analyzer from logs and/or trace files in a Data Retrieval Session can be properly parsed by the PEF Runtime and thereafter displayed in a Message Analyzer viewer.

- **POM** — a binary representation of a set of OPN text files in the form of a decorated syntax tree. These descriptions are utilized by the PEF Runtime to parse messages whenever you run a Live Trace Session, if you load an unparsed trace file in .matu format, if you load a log file such as an ETL, or whenever you reparse a trace file.

- **PEF Runtime** — accepts messages from various components, such as network drivers, ETW instrumented message providers, and logs, and processes them by using the parsing information (compiled protocol descriptions) that reside in the POM. The Runtime component also provides an API that enables Message Analyzer and PowerShell to interface with it. Message Analyzer relies upon the Runtime to capture and parse messages and to provide those messages through its API so that Message Analyzer can access and

display them in selected data viewers.

The PEF Runtime is of central importance to Message Analyzer in performing the following tasks:

- Listening for message packets from network driver interfaces, input adapters, and other components that are instrumented as ETW providers.
- Querying the POM to determine if OPN protocol descriptions exist that correspond to retrieved message packets.
- Constructing OPN representations of retrieved packets, providing that corresponding OPN protocol message descriptions were written and compiled.
- Dispatching the OPN packet representations to internal "endpoints" that are monitored by POM "listeners", or "actors", which in turn decode the packets and pass them to higher endpoints up the processing chain, repeating this process until all packets in the message stack are decoded.
- Enabling Message Analyzer to access the decoded messages through the Runtime API and to display them in a data viewer such as the **Analysis Grid**.
- Allowing PowerShell to access messaging functions through the Runtime API.

- **PEF Driver-Providers** — provide the network interfaces for capturing events and messages that are passed to the Runtime parsing engine. For example, the **Microsoft-PEF-NDIS-PacketCapture** provider captures data on the wire starting at the Data Link Layer; the **Microsoft-PEF-WFP-MessageProvider** captures data above the Network/IP Layer; and the **Microsoft-PEF-WebProxy** provider captures HTTP client browser traffic, unencrypted HTTPS, and other messages at the Application layer.

> **NOTE**
> The **Microsoft-Windows-NDIS-PacketCapture** provider also captures messages at the Data Link Layer and above; however, this provider also has remote capabilities that you can employ in certain scenarios, as described in Built-In Trace Scenarios. In addition, this provider is available on computers running the Windows 8.1, Windows Server 2012 R2, Windows 10, or later operating systems only.

All PEF drivers are instrumented for Event Tracing for Windows (ETW) so they can take advantage of the ETW infrastructure and deliver both events and captured network traffic. In turn, the events and network messages are passed to the Runtime parsing engine and thereafter Message Analyzer can display them.

> **NOTE**
> The PEF Runtime can also parse messages from system ETW providers that exist on your computer, as long as Message Analyzer was able to retrieve a manifest for them during installation. After Message Analyzer successfully finds and stores the manifest for a system ETW provider, you can capture the provider's events during a Live Trace Session. However, you must first select the ETW provider you want to use in the ***Add Providers*** drop-down list on the **ETW Providers** toolbar of the **New Session** dialog during Live Trace Session configuration, as described in Adding a System ETW Provider.

More Information
To learn more about PEF providers and their features, see PEF Message Providers.

- **Input Adapters** — provide the interfaces that define entry points or "chokepoints" into the PEF Runtime for various Import Entities, in message file formats such as .etl, .cap, .log, .matu, and .matp.

PEF architecture also contains other components, such as a POM Adapter that provides importing and exporting facilities; Simulation, which enables modeling of protocol test suites; and technical document (TD) generation, which produces documentation stubs and other artifacts for writers. These components are mentioned here

because they interact with OPN protocol descriptions as part of PEF architecture, but are not directly related to Message Analyzer functions, with the exception of certain POM adapters.

Message Analyzer Integration into the PEF Architecture

The diagram that follows shows how Message Analyzer fits into the PEF architecture.

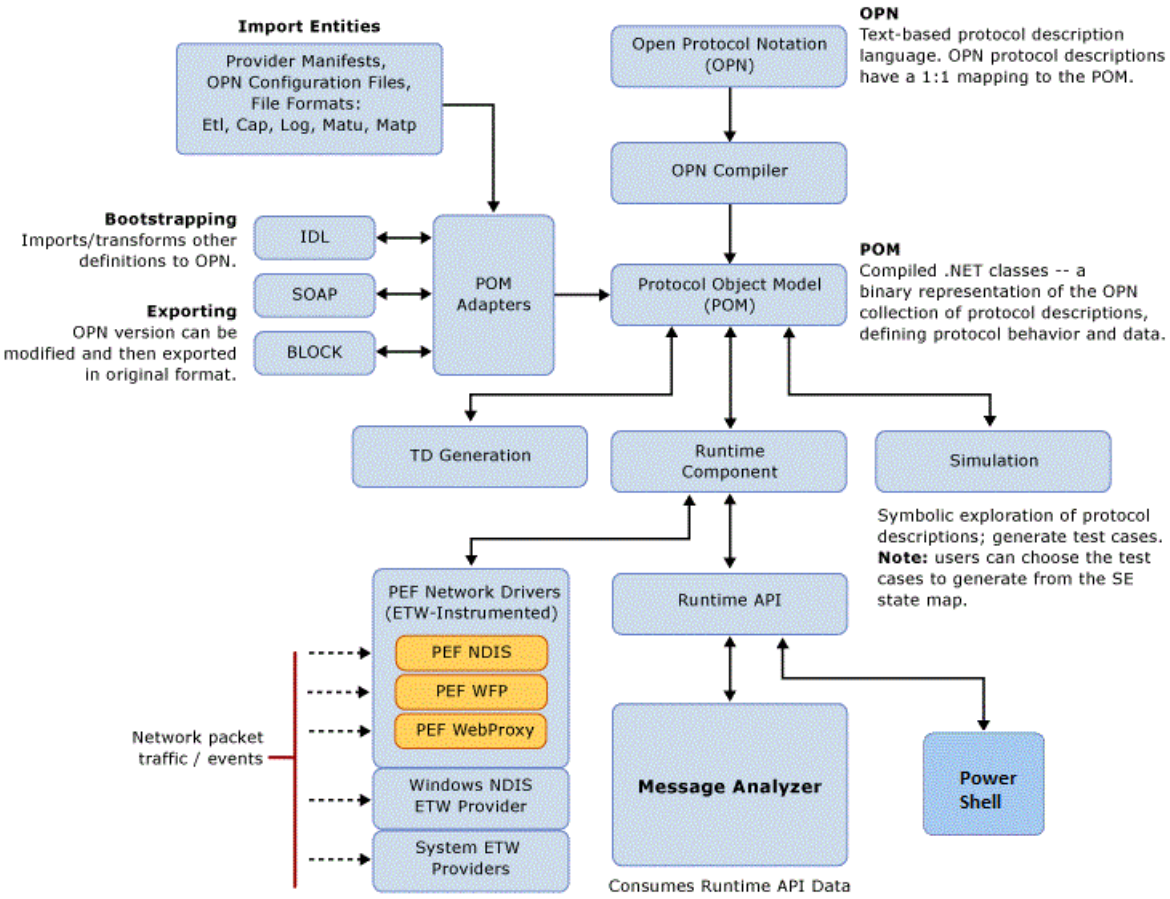

Figure 16: PEF component architecture

More Information

To learn more about OPN programming, see the OPN Programming Guide, which is currently available from the Microsoft download site. In the future, an *OPN SDK* may be available on MSDN to include tutorials, walkthroughs, standard library, language, and other managed reference documentation, depending on demand.

To learn more about viewing the OPN definition for any protocol or module that Message Analyzer parses, see Viewing OPN Source Code.

To review an OPN walkthrough for two TCP Pattern Expressions that are provided by default with Message Analyzer, see Understanding Message Pattern Matching.

To learn more about how to create an OPN configuration file that parses a custom text *.log file, see Parsing Input Text Log Files.

ETW Framework Conceptual Tutorial

11 minutes to read

This tutorial provides conceptual overviews of the underlying Event Tracing for Windows (ETW) framework upon which message/event capturing is based in Message Analyzer.

Architectural Overview

As implemented in the Windows 7 and later operating systems, ETW is a high-speed tracing facility that uses kernel buffering and logging to provide a tracing mechanism for events that are raised by both user-mode applications and kernel-mode device drivers. These events are traced and logged via an ETW Session. The topics that follow provide overviews of the ETW framework and the architecture in which its components exist.

Event Instrumentation

To enable a software component to report critical errors and other important events, it can be instrumented for ETW. If a software component is configured with event instrumentation for key errors and execution states, it can help you do the following:

- Reduce debugging time.
- Resolve performance problems.
- Monitor low-resource conditions or failures in software and hardware components.
- Pinpoint poorly performing components and services.
- Identify bottlenecks.

For a software component to report events, it must register with ETW as an event provider. To facilitate event reporting functionality, a provider configuration must be specified to define the event descriptions, data, and format. Thereafter, when the software component encounters an error condition or other important execution state that has been instrumented for ETW, the provider raises corresponding events. These events need to be recorded somewhere for further processing. In the ETW framework, the component to which provider events are initially written is called an ETW Session. In turn, an ETW Session delivers the event data live to a consumer or logs it for later processing and analysis.

ETW has a unified API that combines the processes of logging and writing trace events to consumers in a single convenient mechanism that easily accommodates event providers. In ETW, event trace sessions are not statically tied to providers, but rather, exist in different spaces to enable more dynamic and flexible trace and event management.

An ETW Controller starts and stops ETW Sessions and dynamically enables providers. This means that the ETW Controller can enable a group of providers for a session, disable some at a certain point, and then enable others later on. Also, because providers are separate from ETW Sessions that operate in the kernel, providers typically have no knowledge of the session to which their events are being logged. An important advantage of separating trace sessions and providers is that it makes tracing immune to application crashes and hangs, such that any events logged by a provider before a crash will be in kernel memory or in a trace file already. This ensures that events are not lost, which is very useful when you are debugging crashed applications.

The high-level components of the ETW session architecture are illustrated in the following figure. The functions of these components, along with the role of event definitions and manifests, are described thereafter.

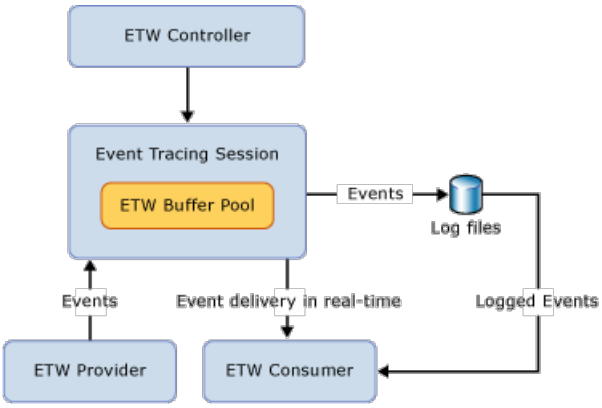

Figure 17: ETW session architecture

More Information

To learn more about how to configure certain parameters of an ETW session from Message Analyzer, see Specifying Advanced ETW Session Configuration Settings.

Event Definition

Trace events are written to the buffer configuration of an ETW Session. The contents of a trace event includes an event header and data that is defined and written by an event provider to describe the current state of an application or process. Every trace event is stamped with a provider ID and is assigned a structure called the Event Descriptor. Both of these data elements exist in the event header. The Event Descriptor defines standard event information with members such as the event Id, Version, Channel (target audience specifier), Level string, Opcode string, Task identifier, and a Keyword assignment value. Keywords are of noteworthy importance given that each one consists of a unique numeric value that distinguishes a different event generated by an ETW Provider, which in turn can be selectively passed to an ETW consumer, as described in System ETW Provider Event Configuration ahead.

The following type definition shows the above event information as fields in an EVENT_DESCRIPTOR structure:

```
typedef struct EVENT_DESCRIPTOR {
    USHORT Id;
    UCHAR Version;
    UCHAR Channel;
    UCHAR Level;
    UCHAR Opcode;
    USHORT Task;
    ULONGLONG Keyword;
} EVENT_DESCRIPTOR, *PEVENT_DESCRIPTOR;
```

The Event Descriptor entries are initially specified in an event manifest that is usually written when a software component is instrumented for ETW. The headers for the trace events of an instrumented component are generated thereafter from the event manifest.

> **NOTE**
>
> When a trace event is logged to an ETW Session, ETW adds other data to the event header, which includes a timestamp, process and thread ID, processor number, and CPU usage data of the logging thread. This data is passed to the ETW Consumer along with the event descriptor information given by the provider. This additional data can be invaluable in trace analysis.

Event Manifest

An event manifest describes the ETW Provider and the data format in which its events are written so that ETW Consumers can process such event data. An event manifest is written in XML, in which the corresponding Event

Descriptor information is specified in XML tags. Event manifests are included for all ETW Providers that are registered on your system.

Event manifests are usually written by developers when designing instrumentation for a software component. The Event Descriptor data items are specified in the event manifest, along with other event metadata and user information. Metadata fields are combined for each event and defined in an <Event> tag that is uniquely associated with an event ID. Event layouts are specified with a <Template> tag which describes user-specified context data for each event. The layout can specify data fields such as strings and integers, or other more complex data structures. Template information does not have to be specified for all events; however, if unspecified for a particular event, it will have no user context data. The following is a fragment of a hypothetical XML event manifest for an ETW Provider:

```
<provider name="Microsoft-Windows-Kernel-SomeComponent"
    guid="{70eb4f03-c1de-4f73-a051-33d13d5413bd}"
    symbol="SomeComponentProvGuid"
    resourceFileName="%SystemRoot%\System32\compapi32.dll"
    messageFileName="%SystemRoot%\System32\compapi32.dll">

<channels>
    <channel name="Microsoft-Windows-Kernel-SomeComponent/Analytic"
        chid="ComponentEvents" symbol="COMP_Events" type="Analytic"
        isolation="System">Contains events for SomeComponent.</channel>
</channels>

<opcodes>
    <opcode value="32" name="CreateComp" symbol="" />
    ...
</opcodes>

<keywords>
    <keyword name="CompCreate" symbol="" mask="0x1000" />
    ...
</keywords>

<templates>
    <template tid="tid_CompCreate">
      <data name="BaseObject" inType="win:Pointer"
         outType="win:HexInt64" />
      <data name="KeyObject" inType="win:Pointer"
          outType="win:HexInt64" />
      <data name="Status" inType="win:UInt32"
          outType="win:HexInt32" />
      <data name="RelativeName" inType="win:UnicodeString"
          outType="xs:string" />
    </template>
  </templates>

<events>
    <event value="1" symbol="ETW_COMPONENT_EVENT_CREATE_OBJECT"
        template="tid_CompCreate" opcode="CreateComp"
        channel="ComponentEvents" level="win:Informational"
        keywords="CompCreate"
        message="$(string.component.compcreate)"/>
        ...
      </events>

</provider>
```

ETW Provider

An ETW Provider is the logical entity that raises events and writes them to an ETW Session. When a software component is being instrumented for ETW, an ETW Provider is created to specify the events it writes, which includes the definition of associated Event Descriptors and the maximum size of each event. The ETW Provider

must also contain code that registers the provider with ETW when it is enabled and code that unregisters the provider when its execution is terminated. When the ETW Provider registers, it specifies a provider ID to ETW.

The following code example illustrates a simple ETW provider that writes one event:

```
#include <myevents.h>   // The header is generated from a manifest and
contains the provider Id and EVENT_DESCRIPTOR structure.

REGHANDLE MyProvRegHandle;
ULONG MyInteger;
PWCHAR MyString;               // User-provided data item
ULONG MyStringLength;     // User-provided data item
EVENT_DESCRIPTOR DataDescriptor[2];

// Register the ETW provider:
Status = EventRegister(&MyProviderId,     // ProviderId (GUID)
                NULL,     // Optional callback
                NULL,     // Optional callback context
             &MyProvRegHandle);     // Registration handle

// Construct a DataDescriptor and write an event with MyInteger and MyString:
EventDataDescCreate(&DataDescriptor[0],     // DataDescriptor
                    &MyInteger,     // Pointer to the data
                    sizeof(ULONG));     // Size of data
EventDataDescCreate(&DataDescriptor[1], &MyString, MyStringLength);

// Invoke EventWrite:
Status = EventWrite(MyProvRegHandle,     // Registration handle
                MyEventDescriptor1,     // Header EVENT_DESCRIPTOR type
                DataDescriptor);     // DataDescriptor array

// Unregister the ETW provider:
Status = EventUnregister(MyProvRegHandle);
```

After the provider registers with ETW, an ETW Controller can then enable or disable event tracing in the provider. The provider usually defines its interpretation of being enabled or disabled in code. Generally, an enabled provider generates events, while a disabled provider does not. When the provider raises events, it invokes the ETW logging API to write the events for which the associated software component has been instrumented. The logging API then sends events to a specific ETW Session that is designated by the ETW Controller.

The two types of providers include the classic provider and manifest-based provider. Message Analyzer makes use of both of these provider types, but mostly manifest-based providers. Manifest-based providers define events in a .man file, while classic providers, such as those based on the managed object format (MOF), use a schema to define their events. Manifest-based providers employ the EventRegister method to register the provider and the EventWrite method to write provider events. By using a manifest, an ETW provider can define its events so that an ETW Consumer knows how to process them.

ETW Session

An ETW Session provides an environment that accepts and buffers the events that are written by an ETW Provider. ETW Sessions typically create a trace file for logging the events and can also deliver the events in real-time to consumer applications such as the PEF Runtime, the output data of which is consumed by Message Analyzer. ETW Sessions are allocated a buffer pool to collect event data written by an ETW Provider. A separate write thread is invoked in the ETW Session to flush the buffer data to the ETW log file and ETW Consumer. See Specifying Advanced ETW Session Configuration Settings for information about modifying these settings.

ETW Controller

An ETW Controller is an application that performs the following tasks:

- Defines the size and location of the ETW log file.
- Starts and stops ETW Sessions.
- Enables providers so they can log events to the ETW Session.
- Manages the size of the buffer pool.
- Obtains execution statistics for ETW Sessions.

Session statistics include the number of buffers used, the number of buffers delivered, and the number of events and buffers lost.

ETW Consumer

An ETW Consumer receives events from ETW Sessions in real time or from a log file. The ETW Consumer can select one or more ETW Sessions as a source of events. The ETW Consumer can also request events from multiple ETW Sessions simultaneously, although they will only be delivered in chronological order. When processing events, an ETW Consumer can also specify an event time frame such that only events occurring within a specific window of time will be delivered.

In the case of Message Analyzer, the PEF Runtime is the initial consumer of ETW events that are retrieved by PEF message providers and other system ETW Providers. Message Analyzer then consumes the PEF Runtime data to display log or trace results.

Microsoft PEF Message Providers

Message Analyzer can capture messages through PEF drivers that have been instrumented as event providers for ETW, which includes the **Microsoft-PEF-NDIS-PacketCapture**, **Microsoft-PEF-WFP-MessageProvider**, and **Microsoft-PEF-WebProxy** providers. This instrumentation enables Message Analyzer to take advantage of the ETW infrastructure for collecting data, controlling sessions, configuring buffers, and passing event data to consumers. As a result, PEF message providers can deliver the ETW packets they capture as events from a well-defined and proven tracing environment. Message Analyzer enables you to use these events to debug protocol communications, applications, and processes, in addition to performance analysis.

System ETW Provider Event Configuration

Message Analyzer can also use system ETW Providers to capture specific data from various Windows components or other applications that have been instrumented for ETW. In addition, you can configure system ETW Providers to filter for specific events of these components, as long as the provider specifies a **Keyword** bitmask and/or **Level** configuration. When a developer is coding a provider, he or she has the option to enhance its ETW instrumentation by assigning **Keyword** and **Level** definitions as 8-byte bitmask and 1-byte integer values, respectively. The **Level** value enables the provider to filter events based on their severity or verbosity, such as critical or informational events, while the **Keyword** value enables the provider to filter for events from specific subcomponents that have been instrumented for ETW tracing. When an ETW Controller enables a provider, it sends the provider's **Keyword** and **Level** specifications to the ETW Session so that the provider only captures the events that correspond to user-specified **Keywords**.

In Message Analyzer, you have the option of setting the **Keyword** and **Level** values of system ETW Providers for events that you want to capture, that is, if they provide such configuration options. When you start a trace, this causes the ETW Controller to enable the provider to trace only the events you specified from a particular subcomponent, specific events from several subcomponents, all events from all subcomponents, and so on.

In Message Analyzer, you can access the **Keyword** and **Level** configurations for system ETW Providers from the **ETW Core** tab of the **Advanced Settings** dialog for any message provider that is selected in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog. The **Advanced Settings** dialog for a provider displays when you click the **Configure** link to the right of the provider in the **ETW Providers** list.

NOTE

When you install Message Analyzer, it enumerates all system ETW Providers that are registered on your computer, organizes the providers into a searchable library, and thereafter enables you to access them from the ***Add Providers*** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog during Live Trace Session configuration.

More Information

To learn more about ETW providers, including how to create and register ETW providers and instrumentation manifests, see Creating an ETW Provider on MSDN.

To learn more about generating ETW manifests with various tools, see Generating a Provider Manifest.

To learn more about configuring system ETW Providers, including how to specify event **Keyword** and **Level** settings, see System ETW Provider Event Keyword/Level Settings.

See Also

Specifying Advanced ETW Session Configuration Settings

Message Analyzer Startup Options

5 minutes to read

This section describes several different ways in which you can start Message Analyzer. Note that the first time you start Message Analyzer after installation, you are prompted by the **Welcome to Message Analyzer** dialog to opt-in or opt-out of automatic updates to Message Analyzer user Library asset collections. To understand what these options mean, refer to Syncing Items on First Startup.

The topics ahead describe various ways of starting Message Analyzer along with the startup options that are available. Also, the option for changing the location of Message Analyzer temporary files to a non-system drive is included.

Using Various Methods to Start Message Analyzer

Message Analyzer startup methods consist of the following:

- Click the **Microsoft Message Analyzer** icon in your computer's **Start** menu or on the system Taskbar.
- Double-click a supported Message Analyzer file, such as a .matp, .matu, or .cap file.

  > **NOTE**
  > When you double-click a supported trace or log file to start Message Analyzer, or if you use the right-click method indicated next, the file's message data is automatically loaded into the **Analysis Grid** viewer by default.

- Right-click a supported Message Analyzer file and select the **Open** item from the context menu. Note that in some cases, you might need to select the **Open with Message Analyzer** item in the context menu, for example, with a .cap or .etl file.
- Specify a startup command string at the command line prompt, as described ahead.

> **NOTE**
> For more information about supported Message Analyzer files, see Locating Supported Input Data File Types.

Starting Message Analyzer From the Command Line

When you start Message Analyzer from the command line, you can do the following:

- Launch Message Analyzer to the **Start Page**, without opening a trace or log file.
- Start Message Analyzer while opening a specified trace or log file, which displays data in the **Analysis Grid** viewer by default.
- Start Message Analyzer while specifying various command line options associated with locating, logging, and compiling any custom OPN parsers that you have created for Message Analyzer use.

The basic syntax for starting Message Analyzer from the command line is the following:

```
MessageAnalyzer.exe [TraceFile|LogFile] [Options]
```

For example, you can open Message Analyzer from the command line without specifying a trace file, log file, or any options, as follows: `"C:\Program Files\Microsoft Message Analyzer\MessageAnalyzer.exe"`

You can also specify a trace or log to open, as indicated in the following example, but it must be in one of the supported file formats, as described in Locating Supported Input Data File Types:

```
“C:\Program Files\Microsoft Message Analyzer\MessageAnalyzer.exe” <TraceFile.matp>
```

Using the Startup Switches

The command line switches that are available for starting Message Analyzer at a command prompt are described below. To view these options at the command line, specify the following help switch:

```
"C:\Program Files\Microsoft Message Analyzer\MessageAnalyzer.exe" /?
```

> **IMPORTANT**
> With exception of the **TraceFile** argument and the command line usage switch, the following list items enable you to locate, compile, log, and debug custom OPN parsers that you have created for use with Message Analyzer.

- **TraceFile** — specifies the name and path to the trace or log file to open when Message Analyzer starts.
- **/?** — displays the command line usage instructions described here.
- **/OPNLoadPathOnly=<path>** — specifies the only path from which you will load custom OPN parsers. If this switch is specified, the default OPN load path in Message Analyzer is ignored. Note that this switch cannot be used with the `/OPNLoadPathMerge` option.
- **/OPNLoadPathMerge=<path>** — adds the specified path to the list of paths from which you will load custom OPN parsers. Note that this switch cannot be used with the `/OPNLoadPathOnly` option.

  > **IMPORTANT**
  > If you create a new path from which to load custom OPN parsers and you place a native Message Analyzer parser in this path, for example, a user-extended native parser with the same name, Message Analyzer will start improperly.

- **/EnableTracing <flag>** — specifies whether tracing is enabled. When set to True, tracing is enabled; when set to False, tracing is disabled (the default value if this switch is not specified).
- **/LogMode=<log mode>** — specifies the output of the compilation log for debugging OPN. The log mode enables you to redirect log messages in the following ways:
  - **/LogMode=File** — saves log messages to a log file. If you do not specify a <filename> value with the `/LogPath` switch, the log file will be created as %LOCALAPPDATA%\Microsoft\MessageAnalyzer\Parsers.log.
  - **/LogMode=Console** — redirects log messages to the console.
  - **/LogMode=ETW** — redirects log messages to the Windows Event Trace system.
  - **/LogMode=Default** — redirects log messages to the Win32 debug output.
- **/LogPath=<file name>** — creates a log file with the specified file name. This switch can be used only with the `/LogMode-File` option.
- **/CachePath=<path>** — specifies a path to an alternate cache folder where custom OPN parser model and codec files will be written and read, so that the default compilation cache for OPN parsers provided with Message Analyzer is not overwritten. Message Analyzer will then point to this folder at startup.
- **/GenerateOPNSymbols** — generates PDB (Program Database) files during OPN compilation for instrumentation and debugging. When a .pdb file is generated, you can attach Message Analyzer to a debugger to debug OPN-generated C# code.

- **/OptimizeOPN=<flag>** — specifies whether or not OPN compilation should be optimized; for example:
  - **/OptimizeOPN=true** — optimizes OPN compilation. True is the default value, if this switch is unspecified.
  - **/OptimizeOPN=false** — does not optimize OPN compilation. Makes it easier for OPN debugging, which is more difficult to perform when OPN code is optimized. Note that if you specify this flag, Message Analyzer will have to perform recompilation at startup time.

Changing the Message Analyzer Temporary Files Location

Message Analyzer enables you to specify a location on a non-system drive for Message Analyzer temporary files. You might do this for the following reasons:

- **Better management of disk space** — avoids excessive use of memory space from Message Analyzer writing to your c:\ drive over a period of time.
- **Better performance** — Message Analyzer has better performance when writing to a non-system drive.
- **Improve performance on Windows 8 and later** — Message Analyzer performs better on the Windows 8 operating system and later if the destination disk is formatted to the Resilient File System (ReFS), instead of the enforced NTFS format of your default system drive.

Although the main application that benefits from reconfiguration of the temporary files location is usually server components, you still have the option change the location. If you want to do this you will need to open the MessageAnalyzer.exe.config file from the following location (or the installation folder location that you specified when installing Message Analyzer), and modify it as indicated below. Note that you could open this XML file in the Visual Studio IDE for the convenience of auto-formatting:

`C:\Program Files\Microsoft Message Analyzer\MessageAnalyzer.exe.config`

Under the `<configuration/>` node in the MessageAnalyzer.exe.config XML file contents, enter the following and specify an appropriate value for the drive location.

```
<appSettings>
<add key="TempFolderPath" value="<driveLocation>\temp\matemp" />
</appSettings>
```

See Also

Starting Message Analyzer for the First Time

Setting Message Analyzer Global Options

5 minutes to read

Message Analyzer provides a global **Options** dialog that is accessible by clicking the **Options** item in the Message Analyzer **Tools** menu. The **Options** dialog is shown in the figure that follows:

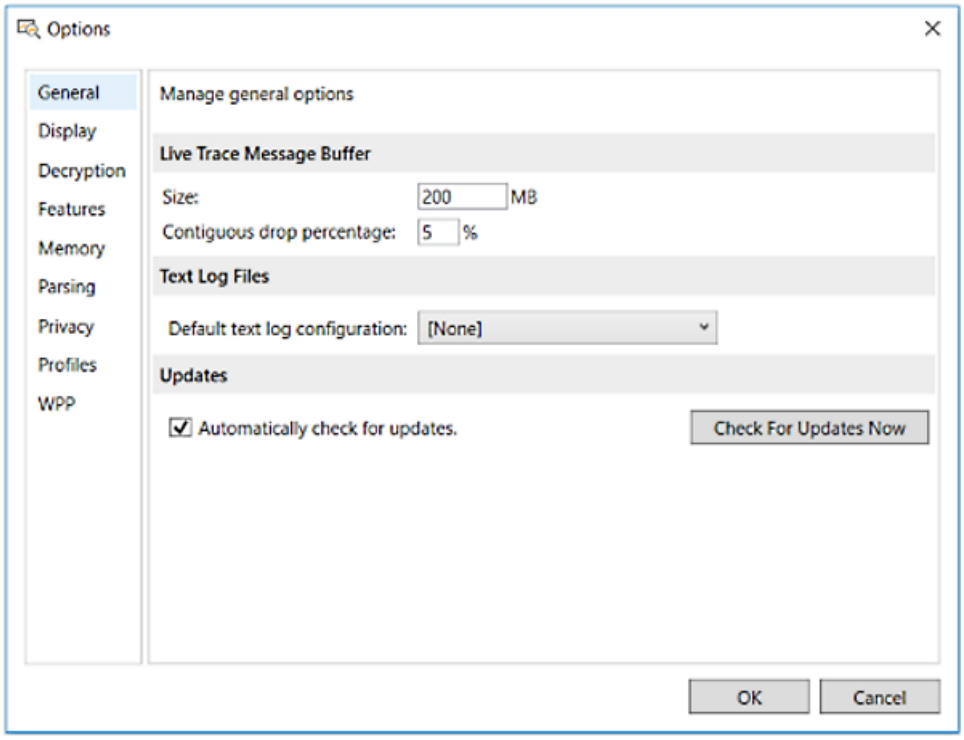

Figure 18: Message Analyzer Options Dialog

In the dialog, you can specify several global options such as default values and selections that can affect Message Analyzer performance, display configurations, feature activation, and parsing functionality. You can set these options at any time; however, you would typically do so prior to starting a Live Trace Session where you capture new data, or prior to a Data Retrieval Session where you load data from saved files into Message Analyzer.

Setting Global Options

The Message Analyzer options that are available from the **Options** dialog on the indicated tabs below consist of the following:

- **General** tab — enables you to specify various default global settings for Message Analyzer, as follows:
  - **Live Trace Message Buffer** pane — provides settings that determine the rate at which packets are dropped when exceeding the buffer limit.

    More Information
    For information about setting the buffer size and count for an ETW session, see Specifying Advanced ETW Session Configuration Settings.

  - **Text Log Files** pane — provides a **Default text log configuration** drop-down list that enables you to select a predefined default or custom configuration file for parsing text logs, as described in Opening Text Log Files. You can specify a chosen value that reflects a text log type that you regularly work with, as described in Parsing Input Text Log Files.

More Information
For more information about creating text log configuration files, see Addendum 1: Configuration Requirements for Parsing Custom Text Logs.

- **Updates** pane — enables you to configure Message Analyzer to **Automatically check for updates** and provide notification when an updated version of Message Analyzer is available for download. Also provides the option to override the automatic setting and check for updates at any time you wish, by clicking the **Check for Updates Now** button.

- **Display** tab — provides the controls that enable you to change the time format and time zone used by Message Analyzer, along with options for the binary value display format used by Message Analyzer.
  - **Time Display** pane — provides settings that enable you to specify the time format used by Message Analyzer. This includes options for **Date and time format**, in addition to a **Reference Time Zone** setting.

    More Information
    For more information, see Configuring Time Format Settings.

  - **Binary Values** — by selecting one of the following options on the **Display** tab of the **Options** dialog, you determine the default format in which Message Analyzer displays binary values, for example, in the **Details Tool Window** and the **Analysis Grid** viewer:
    - **Display as ASCII**
    - **Display as Hex**
    - **Display as Decimal**

  The priThe primary use for this feature is to enable you to choose the data format that displays for the **Payload** field in the **Details** window. For example, by using the default **ASCII** setting, you can cause key information to be exposed in the **Payload** field as text, such as an error string, file path, or other indications that can be useful. Note that you can override the default value in either **Details** or the **Analysis Grid** viewer from the right-click context menus in those locations.

- **Decryption** tab — provides the controls that allow you to import and select server certificates and to specify passwords that are required to enable Message Analyzer to decrypt traffic that is encrypted with the Transport Layer Security (TLS) and Secure Sockets Layer (SSL) security protocols. For example, Message Analyzer can decrypt HTTPS and Remote Desktop Protocol (RDP) messages.

  More Information
  For more information, see Decrypting TLS and SSL Encrypted Data.

- **Features** tab — provides for selection of preview features that you can enable in Message Analyzer. After you select a preview feature, you must restart Message Analyzer for the change to take effect, at which time the feature/s will be available in your asset libraries. A check mark in a preview feature check box will make it available on the next Message Analyzer restart.

> **NOTE**
> Preview features are experimental and may not always work as expected. They are made available so that you can try out new Message Analyzer functionality and provide Feedback on suggested features. Your feedback will help Microsoft make fixes and improvements to Message Analyzer. Note that you can also use the Message Analyzer feedback mechanism that is located in the upper right section of the user interface.

- **Memory** tab — specifies the current memory statistics for Message Analyzer, such as values for **Working Set** and **.NET Current Memory Allocation**, along with the current state (enabled/disabled) of **Server Garbage Collection**. Also contains instructions for how to disable the **Server Garbage Collection** mode to reduce memory consumption, although this could result in lower throughput. The default value for **Server Garbage Collection** is enabled.
- **Parsing** tab — enables you to reparse a set of trace results based on alternate ports that you specify for specific protocols, to accommodate for network traffic that used alternate ports for security purposes. Protocols that are available for alternate port settings are included in the list that follows:

> **NOTE**
> In addition, the **TCP** protocol is included so you can reparse with TCP auto-reassembly disabled to reduce virtual segment message count; you can also reparse with verbose TCP diagnosis messages enabled. The default values for these two features are enabled and disabled, respectively.

- **AllJoyn**
- **HTTP**
- **LDAP**
- **NetFlow/IPFIX**
- **RDP**
- **SMB/SMB2**
- **SSL/TLS**
- **TCP**
- **TDS**
- **TURN**
- **WSRMTCPPort**

After you select an alternate port on which to reparse the messages of a supported protocol, a reload of all data for reparsing will occur immediately after you click the **OK** button to exit the **Options** dialog. Note that you can display the **Parsing** tab of the **Options** dialog by right-clicking a message in the **Analysis Grid** viewer and then selecting the **Parse As...** command in the context menu that appears.

- **Privacy** tab — enables you to opt-in or opt-out of the **Microsoft Message Analyzer Experience Improvement** program. You can also give permission or withhold it from Microsoft to contact you regarding Message Analyzer feature and feedback surveys. Note that a **Feedback** drop-down list is now provided in the upper-right corner of the Message Analyzer user interface.
- **Profiles** pane — enables you to select from among multiple **Profiles** that each provide a preset data viewer and **Layout** configuration that creates a targeted analytical environment for an input file type for which each **Profile** is configured. Also enables you to create your own custom **Profiles**. You can also specify a chosen data viewer as the default **Profile** for all Live Trace Session and Data Retrieval Session results. For example, the **Analysis Grid** viewer is initially set as the Message Analyzer application default.

More Information
For more information, see Working With Message Analyzer Profiles.

- **WPP** pane — enables you to specify symbol resolution information that defines the structure of WPP-generated events, so that Message Analyzer can parse and display such events.

More Information
For more information, see Loading WPP-Generated Events.

Accessibility

3 minutes to read

For compliance with accessibility requirements, Microsoft Message Analyzer includes several keyboard shortcut commands that users can apply along with their mouse equivalents, as described in the table that follows:

Table 3. Accessibility Requirements

TASK	TOPIC REFERENCES	KEYBOARD SHORTCUT	MOUSE EQUIVALENT
Start the configuration for a new session.	Starting a Message Analyzer Session	`Ctrl+N`	Click **New Session** in the Message Analyzer **File** menu or on the **Start Page**.
Stop a running Live Trace Session	See Managing Session Stop, Pause, Resume, and Restart	`Shift + F5`	Click the **Stop** button on the global Message Analyzer toolbar.
Restart a stopped Live Trace Session	See Managing Session Stop, Pause, Resume, and Restart	`F5`	Click the **Restart** button on the global Message Analyzer toolbar.
Display the **Open** dialog to select an input trace file containing data to load into Message Analyzer.	Performing Data Retrieval	`Ctrl+O`	Click the **Open** item in the Message Analyzer **File** menu or click the **Open** button on the global Message Analyzer toolbar.
Trigger a PowerShell cmdlet action in automation scenarios.	Automating Tracing Functions with PowerShell	`Ctrl+Key`	None
Save changes to a trace file.	Saving Message Data	`Ctrl+S`	Click the **Save** item in the Message Analyzer **File** menu, or click the **Save** button on the global Message Analyzer toolbar.
Save a trace file.	Saving Message Data	`Ctrl+Shift+ S`	Click the **Save As** item in the Message Analyzer **File** menu.
Open the **Go To Message** dialog.	Using the Go To Message Feature	`Ctrl+G`	Click the **Go To Message** icon on the **Analysis Grid** viewer toolbar to open the **Go To Message** dialog.
Open the **Find Message** filtering interface.	Using the Find Message Feature	`Ctrl + F`	Click the **Find Message** button on the **Analysis Grid** viewer toolbar.
Apply a view **Filter** to a set of trace results.	Applying and Managing Filters	`Ctrl+Enter` while the cursor is in the Filter Expression text box.	Click the **Apply** button on the appropriate Filter panel of the Filtering toolbar.

TASK	TOPIC REFERENCES	KEYBOARD SHORTCUT	MOUSE EQUIVALENT
Remove a view **Filter** from a set of trace results.	Applying and Managing Filters	`Ctrl+Shift+Enter` while the cursor is in the Filter Expression text box.	Click the **Remove** button on the appropriate Filter panel of the Filtering toolbar.
Open the **New Viewer** drop-down list	Selecting a Session Data Viewer	`Ctrl + Shift + N`	Click the **New Viewer** drop-down list on the global Message Analyzer toolbar.
Select a group of messages in the **Analysis Grid**.	Selecting Messages to Save	`Shift+Down Arrow`	Click a message in the **Analysis Grid** viewer and drag the mouse over the messages you want to select.
Display inline **Details** for an **Analysis Grid** message.	Analysis Grid Viewer	`Ctrl+Enter`	Right-click a message in the **Analysis Grid** viewer and select **Show Details** from the context menu, or double-click the cubed icon next to the message number for which **Details** are required.
Select all hexadecimal values in the **Message Data Tool Window**.	Message Data Tool Window	`Ctrl+A`	Manually select and drag the mouse over all hexadecimal values in the **Message Data Tool Window**.
Display alternate parsing options in the **Options** dialog.	Setting Message Analyzer Global Options	`Ctrl+P`	Right-click a message in the **Analysis Grid** viewer and select **Parse As...** from the context menu.
Copy data for selected fields in the **Message Details Tool Window**. Common for most tools and windows.	Message Details Tool Window	`Ctrl+C`	Select one or more fields in the **Message Details Tool Window**, right-click them, and then select the **Copy Selected Rows** menu item.
Copy a field name for a selected message in the **Message Details Tool Window**. Common for most tools and windows.	Message Details Tool Window	`Ctrl+Alt+C`	Select a field in the **Message Details Tool Window**, right-click it, and then select the **Copy 'Name'** menu item.

Procedures: Quick Start

23 minutes to read

This section contains simple procedures that you can run to start coming up to speed on Message Analyzer features.

Go To Procedures

Proceed to the procedures listed immediately below to learn how to use Message Analyzer features and functions to accomplish basic tasks such as the following:

Displaying Data Quickly From a Saved Trace File
Starting a Live Trace Session with a Built-In Trace Scenario
Starting a Data Retrieval Session
Modifying an Existing Data Retrieval Session
Displaying Different Data Viewers to Change Analysis Perspectives
Creating and Saving a Customized Trace Scenario

> **NOTE**
>
> Although these procedures demonstrate the use of Message Analyzer capabilities in some basic scenarios, they are only a sampling of what you can accomplish with Message Analyzer, given that you can also apply the methodologies described here to many other scenarios. This is also true of other procedural content in this Operating Guide.

> **IMPORTANT**
>
> If you have not logged off from Windows after the first installation of Message Analyzer, please log off and then log back on before performing these procedures. This action ensures that in all subsequent logons following installation, your security token will be updated with the required security credentials from the Message Capture Users Group. Otherwise, you will be unable to capture network traffic in **Trace Scenarios** that use the **Microsoft-PEF-NDIS-PacketCapture** provider, **Microsoft-Windows-NDIS-PacketCapture** provider, or the **Microsoft-PEF-WFP-MessageProvider**, unless you start Message Analyzer with the right-click **Run as administrator** option.

Displaying Data Quickly From a Saved Trace File

The procedure that follows shows you how to use the Message Analyzer **Open** feature to rapidly access and display data from a saved trace or log file.

To quickly open a saved trace file and display its data

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer. Start Message Analyzer with the right-click **Run as Administrator** option if necessary, as described in the previous **Important** note.
2. Click the global Message Analyzer **File** menu and then click **Open** to display the Windows **Open** dialog.

> **TIP**
>
> If you have already opened files with the **Open** command, a **Recent Files** item is available just below the **Open** command on the **File** menu that displays a submenu to the right, from where you can select a file and immediately open it in the default data viewer.

3. In the **Open** dialog that displays, navigate to a saved trace or log file containing the data you want to display and then click the **Open** button to exit the dialog.

   The saved data displays in the default data viewer.

> **TIP**
> You can quickly retrieve data from one or more saved trace files by dragging and dropping them almost anywhere on the Message Analyzer user interface. In drag-and-drop mode, the data retrieved from each file in a selected set is aggregated into a single default session viewer tab. Note that the drag-and-drop function does not work if you are running Message Analyzer as an Administrator, due to varying security contexts that can occur between applications.

You can also drag and drop *.log files to display their data. However, instead of the data immediately displaying in the default data viewer, the **New Session** dialog opens to the Data Retrieval Session input configuration, with the log file/s that you selected as the data source/s for the session. This gives you the opportunity to specify a **Text Log Configuration** file, which is required for parsing *.log files (unless you have already specified a default configuration file on the **General** tab of the **Options** dialog to use for all *.log files, in which case Message Analyzer immediately begins loading the data).

Note that you can specify other session input configurations that include a **Time Filter**, **Session Filter**, or **Parsing Level**, to define the scope of messages to be retrieved. You can also set the **Truncated Parsing** mode, add more files to the session as data sources, and specify the data viewer you want to use.

More Information
To learn more about additional configuration capabilities for a Data Retrieval Session, see Configuring a Data Retrieval Session.

Starting a Live Trace Session with a Built-In Trace Scenario

The procedure that follows shows you how to select the built-in **Loopback and Unencrypted IPSEC Trace Scenario** that uses the **Microsoft-PEF-WFP-MessageProvider** to focus your live data capture above the Network Layer, while minimizing lower-level network traffic. Although this scenario enables you to capture loopback and unencrypted IPSec traffic, this is not the focus of this example.

> **TIP**
> To start a Live Trace Session immediately with no further configuration, you can simply click the **Loopback and Unencrypted IPSEC Trace Scenario** in the **Favorite Scenarios** list on the **Start Page**, where this scenario is accessible by default. You can also locate it in the submenu of the **Favorite Scenarios** item in the Message Analyzer **File** menu to quickly start a Live Trace Session with this scenario. Note that you can set any **Trace Scenario** to Favorite status at your discretion and it will appear in these locations to enable quick session startup.

To start a Live Trace Session with the Loopback and Unencrypted IPSEC trace scenario

1. Launch Message Analyzer as specified in the previous procedure.
2. Click the global Message Analyzer **File** menu and then click the **New Session** item to display the **New Session** dialog.
3. Under **Add Data Source** in the **New Session** dialog, click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features contained in the **New Session** dialog.
4. Select the **Loopback and Unencrypted IPSEC** item in the **Select Scenario** drop-down menu on the **ETW Providers** toolbar of the **Live Trace** tab in the **New Session** dialog.

   The **Microsoft-PEF-WFP-MessageProvider** is added to the **ETW Providers** list on the **Live Trace** tab

of the **New Session** dialog.

5. Optionally, select a built-in **Session Filter** from the centralized Filter Expression **Library**, such as `IPv4Address==<192.168.1.1>` to capture messages that are sent to and received from a specific computer only. The IP address in this example filter is a placeholder for an actual IP address that you must provide.

> **NOTE**
> A **Session Filter** enables you to define the scope of the data capture while at the same time improve performance by limiting how much data you collect.

6. Click the **Start With** drop-down arrow and select the data viewer in which to display your trace results, or use the default **Analysis Grid** data viewer setting.
7. Click the **Start** button in the **New Session** dialog to start capturing data in your Live Trace Session.
8. While the Live Trace Session is running, launch a web browser and click some links to navigate to several web locations. Alternatively, you can start some network application.

   Message Analyzer starts to accumulate messages in the data viewer that you specified.
9. Stop the trace at a suitable point by clicking the **Stop** button on the global Message Analyzer toolbar.

   Inspect your trace results in the data viewer that you chose and observe that Message Analyzer has captured a set of messages, including HTTP, as a result of the browser links that you clicked.

More Information
To learn more about how you might analyze HTTP and TCP message data, see the following topics for some examples of how to apply HTTP and TCP view **Filters** in an Analysis Session:
To apply an HTTP view Filter to Loopback and Unencrypted IPSEC trace results and examine all HTTP-related messages
To apply TCP view Filters to Loopback and Unencrypted IPSEC trace results and expose TCP diagnostics

To learn more about the configuration capabilities that are available for a Live Trace Session, see Configuring a Live Trace Session.

Caution
Be aware that if you let a trace session run for an extended period, it can consume a large amount of memory.

Starting a Data Retrieval Session

The procedure in this section shows you how to open the input configuration for a Message Analyzer Data Retrieval Session, from where you can specify one or more saved files that contain the message data you want to load and display in the **Analysis Grid** viewer. The option to create a filtered view of the loaded data with the use of a **Session Filter** is also described.

To use a Data Retrieval Session to load saved trace data into Message Analyzer

1. Launch Message Analyzer as indicated in earlier procedures.
2. On the **Start Page**, click the **New Session** button to display the **New Session** dialog.
3. Under **Add Data Source** in the **New Session** dialog, click the **Files** button to display the **Files** tab along with the associated session configuration features that the dialog contains.
4. On the **Files** tab, click the **Add Files** button to display the **Open** dialog, from where you can navigate to the trace files that contain the data you want to load into Message Analyzer.
5. In the **Open** dialog, select the file/s that contain the data you want to retrieve, then click the **Open** button

to exit the dialog.

Message Analyzer displays the files you selected in a list on the **Files** tab that includes columns of data such as **Name**, file **Size**, **File Type**, **Message Count**, **Start Time**, **End Time**, and **Text Log Configuration**.

> **NOTE**
> The data from the files that display in this list is not yet loaded into Message Analyzer. At this point, the files are simply the target data sources from which data will be loaded after you click the **Start** button in the **New Session** dialog.

6. In the files list, ensure that there is a check mark in the check box next to the file/s containing the data you want to load into Message Analyzer. Note that you can select or unselect files in the list to create specific combinations of data sources from which to load data.
7. In the **Start With** drop-down menu of the **New Session** dialog, select a data viewer in which to display the results of your Data Retrieval Session; otherwise, the default data viewer setting will be used.

> **TIP**
> You have the option to change the default data viewer from the **Profiles** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. Note that Message Analyzer ships with the **Analysis Grid** viewer as the default setting, as specified in the **Default Profile** pane of the **Profiles** tab.

8. Optionally, select a **Session Filter** from the Filter Expression **Library** in the **New Session** dialog, or configure a **Time Filter** from the **Files** tab in the dialog to define the scope of data retrieval or to narrow the window of time in which to view data, respectively.

More Information
To learn more about applying a **Session Filter** and/or a **Time Filter**, see Selecting Data to Retrieve.

---

9. Click the **Start** button in the **New Session** dialog to begin loading the data into Message Analyzer.

   After the data is loaded, it displays in the default or selected data viewer.

More Information
To learn more about how to manipulate and analyze saved trace data that you have loaded into the Message Analyzer **Analysis Grid** viewer, see the following sections:

Analysis Grid Viewer
Common Data Viewer Features
Tool Windows
Filtering Live Trace Session Results

Modifying an Existing Data Retrieval Session

If you want to modify an existing Data Retrieval Session so that you can load additional data from one or more files, for example saved traces and logs, or if you want to specify other configurations that include a **Session Filter**, perform the steps of the following procedure.

To modify an existing Data Retrieval Session

1. Load data into Message Analyzer through a Data Retrieval Session, as described in the previous procedure.

2. On the global Message Analyzer toolbar, click the **Edit Session** button to return to the original configuration of the currently in-focus Data Retrieval Session. If you have data for more than one session displaying, ensure that you select a viewer tab for the Data Retrieval Session that you want to modify before you click **Edit Session**.

3. On the **Files** tab of the **New Session** dialog, click **Add Files** to add one or more saved trace files to the files list and then select the check box next to each file containing the data you want to load into Message Analyzer.

> **IMPORTANT**
>
> When you click the **Edit Session** button on the global Message Analyzer toolbar, by default your Data Retrieval Session opens in **Restricted Edit** mode. This mode enables you to add saved files as new data sources, but it disables other configuration capabilities, such as setting a **Time Filter**, choosing a **Session Filter**, or setting the **Parsing Level**.
>
> The advantage of the **Restricted Edit** mode is that you can add the new data files without triggering a reload of all data and incurring a performance hit. However, if you want to enable the indicated configuration features to specify changes, you can select the **Full Edit** mode, although you should be aware that a reload of all data will be required if you **Apply** the changes.

4. When you are done with reconfiguring the session, click the **Apply** button to load and display the new data in the **Analysis Grid** viewer.

> **NOTE**
>
> When you load data from additional files in an edited Data Retrieval Session, the messages from these files are interlaced with the existing messages in the **Analysis Grid** viewer in chronological order.

Configuring a Session Filter When loading data from saved files into Message Analyzer, you can select a built-in Filter Expression from the **Library** drop-down list above the **Session Filter** text box, or you can manually configure one in the same text box. This results in filtering the input messages to specific criteria.

For example, you might add a simple expression such as `*Port != IANA.Port.LDAP` from the **Library** drop-down list to remove LDAP traffic on TCP and UDP transports, if you don't want to view this data. Note that if you manually configure a Filter Expression and it is invalid, you may receive a **Compile query error** message after you click the **Apply** button in the **New Session** dialog.

> **TIP**
>
> After loading a collection of messages from specified files and displaying the data in a selected viewer, you have the option to add a built-in or manually-configured view **Filter** to further isolate specific data of interest. The centralized Filter Expression **Library** for selecting built-in **Filters** is available on the Filtering toolbar that appears above every session viewer.

More Information

To learn more about how to manually configure your own Filter Expressions, see Writing Filter Expressions.

Displaying Different Data Viewers to Change Analysis Perspectives

The procedure that follows runs a the **Loopback and Unencrypted IPSEC Trace Scenario** in a Live Trace Session and then loads a message collection to create initial data views in separate **Analysis Grid** viewer tabs. You will then select several different data viewers that provide high-level data summaries and statistics, some in graphic formats.

To display different data viewers

1. Follow the guidelines of the second procedure in this section to start a Live Trace Session with the **Loopback and Unencrypted IPSEC Trace Scenario**.
2. Capture SMB traffic by performing file access activities while your Live Trace Session is running.

   Note that the **Microsoft-PEF-WFP-MessageProvider** in the **Loopback and Unencrypted IPSEC Trace Scenario** captures data above the IP/Network Layer, which makes it a good choice for capturing SMB traffic at the Application Layer while minimizing lower layer noise.

> **TIP**
> You might consider using the **SMB2 Client Full Payloads Trace Scenario** to capture SMB traffic unencrypted.

3. Stop the Live Trace Session at a suitable point by clicking the **Stop** button on the global Message Analyzer toolbar.
4. Load messages from one or more saved trace files (preferably related SMB data, if you have it) into Message Analyzer through a Data Retrieval Session by following the guidelines of the third procedure in this section.

   The trace results and loaded data display in separate **Analysis Grid** viewer tabs, assuming that you specified the **Analysis Grid** as your data viewer in the **New Session** dialog for your Live Trace Session and Data Retrieval Session configurations.
5. If the **Session Explorer Tool Window** is hidden, click the global Message Analyzer **Tools** menu, select the **Windows** submenu, and then click the **Session Explorer** item to restore it to the default location.
6. To create different views of the live trace results data (in addition to the existing **Analysis Grid** viewer instance), right-click the live trace session node in **Session Explorer**, highlight **New Viewer**, select **Chart** in the drop-down list, and then select the **SMB File Stats** view **Layout** so that you can analyze file access duration times, bytes transferred, and data transmission rate for each file name in a set of trace results.

   You might also display the **SMB Top Commands** view **Layout** to obtain a high-level overview of message volume per file access operation, as an indication of the SMB operations consuming the most bandwidth.

   Note that each viewer in the same session is identified by a unique color code. Therefore, the color code for each viewer in the Live Trace Session will be the same, while at the same time different from that of the viewer for the Data Retrieval Session in which you loaded data from saved files into Message Analyzer.

More Information
To learn more about the **Layouts** that are available for the **Chart** viewer, see the Chart Viewer Layouts topic.

---

7. Double-click any bar element in the **SMB Top Commands Bar** graph visualizer component **Layout** to display the corresponding SMB messages in a separate **Analysis Grid** session viewer tab for further analysis.
8. Repeat step 6 and select the **SMB Reads and Writes Layout** for the **Chart** viewer to display SMB statistics in the **Timeline** visualizer component.

> **IMPORTANT**
> These viewers will display data only if SMB, SMB2, or SMB3 protocol packets were captured in the Live Trace Session that you ran.

9. Right-click the node for the Data Retrieval Session in **Session Explorer**, highlight **New Viewer**, select the **Grouping** viewer, and then select the **Process Name and Conversations** view **Layout** in the drop-down list that appears.

   Note that the hierarchy of message groups that display for this **Layout** isolate trace data into nested groups consisting of **ProcessName**, **ProcessId**, **Network** (IP conversations), and **Transport** (ports that carried the IP conversations), to enable a focused analytical perspective. By clicking any Group in the hierarchy, you can drive the display of Group messages into the **Analysis Grid** viewer for inspection of associated message data. If you are a Microsoft Network Monitor user, you might notice that this particular **Layout** is similar to the **Network Conversation Tree**.

10. While the **Grouping** viewer has focus, click the **Layout** drop-down list on the **Grouping** viewer toolbar and select the **File Sharing SMB/SMB2** view **Layout** from the **File Sharing** category, to isolate the data by SMB **SessionIdName**, **TreeIdName**, and **FileName**.

    This will enable you to drill down into the hierarchically organized Group display to analyze the messages associated with each file, as nested under the SMB session and tree IDs. for example to search for diagnostic errors. Note that when you select any Group node under the **Group Values** column of the **Grouping** viewer, you will display the messages associated with the selected Group node in the **Analysis Grid** viewer for further analysis, provided that the **Filtering Mode** is enabled on the **Grouping** viewer toolbar. If the **Selection Mode** is enabled, selecting a Group will only highlight the associated messages in the **Analysis Grid** viewer.

More Information
To learn more about the **Grouping** viewer, including the **Filtering Mode** and **Selection Mode** of operation, refer to the Grouping Viewer topic.

---

11. Next, right-click the node for the Data Retrieval Session in **Session Explorer**, highlight **New Viewer**, and then select **Pattern Match** to display the **Pattern Match** viewer.

    To start the **Pattern** matching process, specify a built-in **Pattern** expression in the **AVAILABLE PATTERNS** list of the **Pattern Match** viewer, by selecting the check box of a chosen **Pattern** expression. For example, you might select the **TCP Retransmit Pairs** or **TCP Three-Way Handshake Pattern** expression to identify these types of sequential pattern matches across the retrieved message set, possibly to expose network or connectivity issues, respectively.

More Information
To learn more about **Pattern** matching, refer to the Pattern Match Viewer topic.

---

12. To quickly vary your analysis perspectives, poll through the various views of data by clicking the viewer nodes under each session name in **Session Explorer**, or select different viewer tabs in the Message Analyzer main analysis surface.

    As you select viewer nodes in **Session Explorer**, the data for each of the selected viewer nodes displays in separate viewer tabs in the main analysis surface. As you poll the data views, you obtain unique perspectives that enhances your analysis of the data.

13. Optionally, to obtain alternate but integrated views of the saved message data, select the **Message Stack Tool Window** from the **Windows** submenu of the global Message Analyzer **Tools** menu, if the window is not already open, to expose the underlying message stack that supported top-level transactions. Also select the **Diagnostics Tool Window** from the same menu location to display summary groups of the different types of diagnosis errors that occurred in the retrieved data.

IMPORTANT

The **Diagnostics** window is currently a preview feature. To use this tool, you must enable it on the **Features** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu, and you must then restart Message Analyzer.

TIP

You can compare Live Trace Session results with related data that is loaded into Message Analyzer from a Data Retrieval Session. This provides a convenient method for analyzing current and historical data side-by-side. To learn how to display data viewer tabs side by side, see Redocking Data Viewers and Tool Windows.

Creating and Saving a Customized Trace Scenario

In the procedure that follows, you will create and save a **Trace Scenario** to serve as a trace template with predefined tracing functionality that you can run on demand. The **Trace Scenario** specified in this simple example enables you to isolate traffic to a specific IP address, where you can use two different methods of filtering to achieve that result.

To create and save a Trace Scenario

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.
2. On the Message Analyzer **Start Page**, click the **New Session** button to display the **New Session** dialog.
3. Under **Add Data Source** in the **New Session** dialog, click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features that the **New Session** dialog provides.
4. Select the **Local Network Interfaces Trace Scenario** in the **Select Scenario** drop-down list on the **ETW Providers** toolbar of the **New Session** dialog.

NOTE

If you are running the Windows 7, Windows 8, or Windows Server 2012 operating system, the **Microsoft-PEF-NDIS-PacketCapture** is added to the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog. Otherwise, for later operating systems, the **Microsoft-Windows-NDIS-PacketCapture** provider (with remote capabilities) is added to the list.

5. In the earlier operating system scenario, click the **Configure** link to the right of the **Microsoft-PEF-NDIS-PacketCapture** provider in the **ETW Providers** list to display the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog. Select the **Provider** tab in the dialog and then specify the configurations that follow:
   - In the **Name** column under **System Network**, expand the **Machine** node, and then under **Adapters**, make sure that the **In** and **Out** check boxes for the Ethernet network adapter are selected. This ensures that the **Trace Scenario** will capture both inbound and outbound traffic on the Ethernet adapter. Deselect these check boxes for all other listed adapters.
   - In the **Fast Filters** pane of the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, click the black arrow next to the **Filter 1** designator in **Group 1** and select the **IPv4Address** option from the drop-down menu that displays.

> **NOTE**
>
> With a low-level IPv4 address **Fast Filter**, the **Trace Scenario** will deliver messages to the PEF Runtime that transited to or from a specified IPv4 address only, as the **Trace Scenario** is running. This avoids the additional parsing that would normally be required if you specify a similar **Session Filter** instead, thereby improving Message Analyzer performance.

- Specify an IPv4 address value in the format *192.168.1.1* in the text box adjacent to the drop-down menu, to isolate traffic to the specified IPv4 address. Make sure to substitute appropriately for the IP address placeholder italics value specified in this example.
- Highlight the row in which the Ethernet adapter exists in the **System Network** tree grid of the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, and then click the **Apply to Highlighted** button in **Group 1**.

  The name of the Ethernet adapter displays as the **Target** of the filter **Group**. Click **OK** to exit.

> **NOTE**
>
> Instead of configuring a **Fast Filter**, you can optionally specify a **Session Filter** such as `IPv4.Address == 192.168.1.1` in the **Session Filter** text box of the **New Session** dialog. However, you should note that a **Session Filter** requires more processing time, as indicated earlier. If you choose to use a **Session Filter**, you can remove the previously set **Fast Filter** configuration.

More Information
To learn more about the filtering configurations that you can specify for the **Microsoft-PEF-NDIS-PacketCapture** provider, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

---

6. In later operating system scenarios, click the **Configure** link to the right of the **Microsoft-Windows-NDIS-PacketCapture** provider in the **ETW Providers** list to display the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog. Select the **Provider** tab in the dialog and then specify an IP Address filtering configuration.

   You can do this by first selecting the Ethernet adapter in the **Adapters** list of the dialog and then specifying an IP address in the **IP Addresses** text box in the **Filters** pane of the dialog. Make sure that the selected Ethernet adapter is the only one with a check mark in the corresponding **Enabled** check box; all others should be removed. Note that you could specify a MAC address for the Ethernet adapter instead of, or in addition to, an IP address.

More Information
To learn more about special filtering configurations that you can specify for the **Microsoft-Windows-NDIS-PacketCapture** provider, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

---

7. Click the **Save Scenario** button on the **ETW Providers** toolbar to display the **Edit Trace Scenario** dialog and then specify values for the **Name**, **Description**, and **Category** fields.

> **NOTE**
>
> When you run your customized **Trace Scenario**, the trace results will display in the default data viewer that is specified in the **Default Profile** pane on the **Profiles** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. If you want to change the default viewer from this location, simply click the **Default Viewer** drop-down list and choose a viewer that you want as the default.

8. When your **Trace Scenario** configuration is complete, click the **Save** button to exit the **Edit Trace Scenario** dialog.

Running the Custom Trace Scenario When you save a customized **Trace Scenario**, it becomes a new item in the **Trace Scenario** user Library in the **Category** that you specified. This **Category** will be located under the **My Items** top-level category of the Library, from where you can select and run it at any time. It also becomes part of the Message Analyzer Sharing Infrastructure, which enables you to mutually share the scenarios in the **Trace Scenario** user Library with others.

> **TIP**
> After you run a custom **Trace Scenario** template from the **New Session** dialog, you have the option to reopen the session configuration by clicking the **Edit Session** button on the global Message Analyzer toolbar. Thereafter, you can reconfigure the **Trace Scenario** as required and save the new template configuration again by clicking **Save Scenario**.

More Information
To learn more about creating **Trace Scenario** templates, see Creating and Managing Custom Trace Scenarios.
To learn more about managing the **Trace Scenarios** Library as part of the Message Analyzer Sharing Infrastructure, see Managing Trace Scenarios.

See Also

Capturing Message Data
Retrieving Message Data
Data Viewers

Starting a Message Analyzer Session

10 minutes to read

The basic unit of workflow and starting point of all Message Analyzer operations is a session. A *session* is the instrument that you will configure and then use to collect data that you want to analyze, whether you capture it live on a network or load it into Message Analyzer from saved message files, logs, or other sources. The user interface (UI) for configuring a new session is designed to provide an optimal workflow experience and enables you to utilize essentially two types of Message Analyzer sessions to acquire input data, as follows:

- **Live Trace Session** — enables you to specify a session configuration that captures message data live from network traffic, system components, or devices.
- **Data Retrieval Session** — enables you to specify a session configuration that acquires input message data from multiple disparate data sources, such as saved files, logs, and others.

Go To Quick Session Startup
To learn how to immediately start a Message Analyzer session that captures live data at the Link Layer with a single click, see the Quick Session Startup topic.

Streamlining Session Input Workflow
The Message Analyzer UI provides a streamlined input workflow architecture that is designed to support expansive data collection capabilities, which includes the future addition of new data sources. Currently, Message Analyzer data collection capabilities include the simultaneous capture of live data from a local host and/or multiple remote target computers, as described in Configuring a Live Trace Session, Configuring a Remote Capture, and Configuring Session Scenarios with Selected Data Sources; in addition to retrieval of data from multiple saved data sources, as described in Locating Supported Input Data File Types and Acquiring Data From Other Input Sources.

Message Analyzer also provides an integrated session configuration workflow that minimizes the number of clicks and dialogs you need to get a Live Trace Session or a Data Retrieval Session running, while making configuration features that are common to both session types readily accessible and easy to configure.

Configuring a New Session: Overview
All of the foregoing capabilities are unified through the use of a common interface that is called the **New Session** dialog. This dialog is accessible by clicking the **New Session** item in the Message Analyzer **File** menu, or by simply clicking the **New Session** button on the Message Analyzer **Start Page** or on the global Message Analyzer toolbar. From the **New Session** dialog, you can choose the type of session you want to start, such as a Live Trace Session or a Data Retrieval Session:

- **Live Trace Session** — you can begin the configuration of a Live Trace Session by clicking the **Live Trace** button under **Add Data Source** in the **New Session** dialog. The configuration features that are unique to this type of session enable you to do the following:
  - Target local and/or remote computers from which to capture live data with the use of the **Edit Target Computers** dialog, which is accessible by clicking the **Edit** button next to the **Target Computers** text box, as described in Configuring a Remote Capture.
  - Choose a built-in **Trace Scenario** from the **Select Scenario** drop-down list that includes one or more message providers that capture network messages through different stack layers, or from system components and other devices, as described in Selecting a Trace Scenario.
  - Add one or more system ETW Providers to any **Trace Scenario** to enhance the scope of data

retrieval, by selecting them from the **Add Providers** drop-down list, as described in Adding a System ETW Provider.

- Specify various filtering configurations for message providers, including **Fast Filters**, **WFP Layer Set** filters in **Trace Scenarios** that use the **Microsoft-PEF-WFP-MessageProvider**; event **Keyword** and error **Level** filters; adapter filters and logically chained **Fast Filter** groups in **Local Network Interfaces** traces; NDIS stack, Hyper-V-Switch extension layer, and other special filters in traces that use the **Microsoft-Windows-NDIS-PacketCapture** provider, for example **Remote Network Interfaces** traces that have remote capture capability; and so on, as described in Configuring a Live Trace Session.
- Specify **ETW Session Configuration** parameters to control performance of the underlying ETW sessions in all traces, as described in Specifying Advanced ETW Session Configuration Settings.

- **Data Retrieval Session** — you can begin the configuration of a Data Retrieval Session by clicking the **Files** button under **Add Data Source** in the **New Session** dialog. The configuration features that are unique to this type of session enable you to do the following:
  - Locate saved files and logs from which to retrieve message data, as described in Locating Supported Input Data File Types.
  - Enable the **Truncated Parsing** mode to initiate a smaller Message Analyzer parsing set that improves performance when loading saved files that contain truncated messages, for example a .cap file, as described in Detecting and Supporting Message Truncation.
  - Decrypt the data in saved files, as described in Decrypting Input Data.
  - Configure a **Time Filter** that specifies a window of time in which to view data, to focus results and improve performance, as described in Applying an Input Time Filter to a Data Retrieval Session.
  - Configure and apply a **Session Filter** to focus results on specific data and improve performance, as described in Applying a Session Filter to a Data Retrieval Session.
  - Select a **Text Log Configuration** file for parsing a custom textual log file, as described in Opening Text Log Files.
  - Load data that is derived from execution of PowerShell scripts, as described in Deriving Input Data with PowerShell Scripts.
  - Load data from Azure storage tables and from text log files that are stored in binary large object (BLOB) containers, as described in Handling Azure Data.
  - Load data from various other input sources such as system Event Logs, SQL databases, WPP-generated events, and OMS log data, as described in Acquiring Data From Other Input Sources.

The configuration features of the **New Session** dialog that are common to both session types enable you to do the following:

- Specify a **Session Filter** that limits the data you capture or retrieve based on specific filtering criteria.
- Select a **Parsing Level** to limit how far up the stack Message Analyzer will parse, for better performance and focused message sets, as described in Setting the Session Parsing Level.
- Choose a data viewer in which to display your trace results by selecting one in the **Start With** drop-down list, or use the default viewer setting, as described in Selecting a Session Data Viewer.
- Provide a **Name** for the new session, as described in Naming a Session.

Using Session Options

The **New Session** item in the Message Analyzer **File** menu provides the following submenu items as options

for creating a session:

- **Blank Session** — creates an unconfigured new session.
- **From Current Session** — creates a new session based on the configuration of a currently open and selected session.

Creating a New Session: Workflow Overview
To familiarize yourself with the workflow you will generally follow when creating a basic session configuration, follow the steps below to create a session that captures data live or retrieves data from saved files and logs:

General Workflow Process

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer. If you have not logged off and back on after first installing Message Analyzer, then restart Message Analyzer and use the right-click **Run as Administrator** option.
2. Click the **New Session** button on the Message Analyzer **Start Page** to display the **New Session** dialog.

   Alternatively, click the **New Session** item in the Message Analyzer **File** menu to display the **New Session** dialog.
3. To start session configuration, do one of the following under **Add Data Source** in the **New Session** dialog:
   - Click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features that it contains in the **New Session** dialog, to begin configuration of a Live Trace Session.
   - Click the **Files** button to display the **Files** tab along with the associated session configuration features that it contains in the **New Session** dialog, to begin configuration of a Data Retrieval Session.
4. If you chose to configure a Live Trace Session, do the following to create a basic configuration; otherwise, proceed to step 5:
   - Use the default **Localhost** setting in the **Target Computers** list on the **Live Trace** tab of the **New Session** dialog, to capture data from the local computer only.
   - Choose a built-in **Trace Scenario** in the **Select Scenario** drop-down list on the **ETW Providers** toolbar of the **Live Trace** tab to define the scope of the data you will capture.
   - Optionally, select a built-in Filter Expression from the **Library** drop-down list on the toolbar above the **Session Filter** text box to create a **Session Filter** that narrows the scope of data capture, so you can focus on a specific type of data in your trace results.
   - Choose a data viewer in which to display the results of your Live Trace Session by selecting it in the **Start With** drop-down list of the **New Session** dialog; or accept the default setting.
   - To review additional configuration options for a Live Trace Session, see Configuring a Live Trace Session.
5. If you chose to configure a Data Retrieval Session, do the following to create a basic configuration:
   - Navigate to the trace or log files that contain the data you want to load into Message Analyzer by clicking the **Add Files** button on the toolbar of the **Files** tab in the **New Session** dialog; this action launches the **Open** dialog.
   - From the **Open** dialog, locate and select the files that contain the data to be loaded into Message Analyzer. Click **Open** to exit the dialog.

The file names appear in the files list on the **Files** tab. Note that this list indicates only the supported files that are targeted for file reads during the data loading process, given that at this point, no data has yet been loaded into Message Analyzer.

- Optionally, select specific files in the files list on the **Files** tab of the **New Session** dialog to create a collection of messages to be loaded from a specified subset of files.
- Optionally, configure a **Time Filter** to define a window of time in which to retrieve data and create focus on a smaller and targeted set of retrieved messages for enhanced analysis.
- Optionally, select a built-in Filter Expression from the **Library** drop-down list on the toolbar above the **Session Filter** text box to create a **Session Filter** that narrows the scope of data retrieval, so you can focus on a specific type of data in your retrieved results.
- To review additional configuration options for a Data Retrieval Session, see Configuring a Data Retrieval Session.

What You Will Learn

In the topics of this section, you will learn in detail how to perform the tasks that are described below, which includes how to: capture data live, retrieve saved data, configure different session scenarios, and edit an existing Data Retrieval Session or Live Trace Session.

In This Section

Capturing Message Data — learn how to use Message Analyzer assets to target live data as an input source in a local or remote capture; configure, run, and edit a Live Trace Session; create your own custom **Trace Scenarios**; and use various methods to perform a live capture.

Retrieving Message Data — learn how to use Message Analyzer to target saved data and other repositories as an input source; configure, run, and edit a Data Retrieval Session; and use various methods to perform data retrieval.

Configuring Session Scenarios with Selected Data Sources — learn how to use Message Analyzer's flexible session framework to create session configurations based on single or multiple data sources. Also review some example scenarios and guidelines for Data Retrieval Sessions and Live Trace Sessions.

Editing Existing Sessions — learn how to reconfigure an existing session to obtain different results.

Go To Procedures

To proceed directly to procedures that demonstrate how to use the network tracing features, see **Procedures: Using the Network Tracing Features**.

To proceed directly to procedures that demonstrate how to browse for saved messages, retrieve selected data from saved files, and view the resulting message set, see **Procedures: Using the Data Retrieval Features**.

Capturing Message Data

5 minutes to read

This section describes how to create and configure a new Live Trace Session with Message Analyzer, so that you can quickly begin capturing data from your system. This section also discusses how to focus your Live Trace Session on capturing messages that contain specific types of data, through the selection or modification of various settings, for example selecting one or more message providers, adding built-in or custom filters, and/or specifying a **Parsing Level** to the capture configuration. Other capabilities enable you to focus on capturing remote traffic on specified hosts and virtual machines (VMs), in addition to decrypting certain types of data.

Go To Session Configuration
Go directly to an overview of Live Trace Session configuration workflow, filtering options, and the features that are available for configuring and starting a new Live Trace Session:
Configuring a Live Trace Session

What You Will Learn
In the topics of this section, you will learn how to accomplish the tasks indicated below.

Targeting Live Data as an Input Source — learn how to target specific types of message data to capture, with the use of built-in (default) **Trace Scenarios** and the message providers they contain. Also learn about accessing the built-in **Trace Scenarios** that Message Analyzer provides, review options for instantly starting a Live Trace Session with a single click where no further configuration is needed, and read an overview about optimizing capture configurations. In the subtopics of this section, you can also review detailed information about the following:

- Microsoft Protocol Engineering Framework (PEF) message providers, along with various types of filters you can use to narrow the focus of live captures.
- The **Microsoft-Windows-NDIS-PacketCapture** provider, with its remote capture capability and unique filtering configurations.

Configuring a Live Trace Session — examine an overview of Live Trace Session configuration workflow, read a **Trace Scenario** configuration overview, and review Live Trace Session configuration features that enable you to perform tasks such as:

- Selecting a built-in **Trace Scenario**.
- Using custom **Trace Scenario** templates.
- Adding system ETW Providers to a Live Trace Session and modifying settings for PEF and other ETW Providers.
- Selecting specific data from a Live Trace Session with the use of a **Session Filter** or by setting a **Parsing Level**.
- Modifying the settings of the underlying ETW session, to which message providers are enabled.
- Using the remote tracing capabilities of Message Analyzer.
- Selecting a data viewer for the trace results.
- Using the Message Analyzer decryption feature.

Creating and Managing Custom Trace Scenarios — learn how to create your own custom **Trace Scenarios**,

save and manage your scenarios, run them on demand, and share them with others.

Performing a Live Capture — learn about the methods you can use to perform a live capture, which includes automation that enables you to instantly start a live capture with a single click, or manually customizing the capture configuration of a Live Trace Session before you run it.

Procedures: Using the Network Tracing Features — perform example procedures that demonstrate various aspects of live capture functionality, as discussed throughout this section.

Using the Default Trace Scenarios

When configuring a Live Trace Session, you are advised that your initial approach should be to use one of the default **Trace Scenarios** that each contain one or more message providers. By learning about the type of data you can capture with the use of default providers, you will understand better how to choose the **Trace Scenario** that you require based on the provider/s that such scenarios contain. To further enhance and optimize a default **Trace Scenario**, you can specify various filtering configurations, as described in this section.

Considering System ETW Providers

Only if using a default **Trace Scenario** fails to isolate the data you wish to capture should you consider customizing a Live Trace Session with one or more system ETW Providers that are installed and registered on your computer. For example, you might end up capturing more data than expected with one of the default **Trace Scenarios** such that packets are being dropped by Message Analyzer. If this is the case, you might consider changing your capture configuration to use one or more selected system ETW Providers instead of a built-in **Trace Scenario** to more finely tune the scope of data capture. You also have the option to *customize* any of the built-in **Trace Scenarios** by either adding selected system ETW Providers to the scenario or by modifying the existing system ETW Provider configuration in the scenario. If you elect to customize using either of these methods, Message Analyzer enables you select specific data to capture by providing facilities to modify the **Keyword** and/or **Level** filtering configuration of any system ETW Provider, as described in System ETW Provider Event Keyword/Level Settings.

> **TIP**
> Modifying ETW Provider filtering configurations is applicable to the PEF providers, as described in the PEF Message Providers section, given that PEF providers are also instrumented for ETW **Keywords** and error **Levels**. In addition, these modifications include configuration of **Keyword** and **Level** filters for the **Windows-NDIS-PacketCapture** provider, as described in the Microsoft-Windows-NDIS-PacketCapture Provider topic. Note that customizing a built-in **Trace Scenario** to optimize your overall capture configuration can also include the application of many types of filters other than **Keyword** and **Level** filtering for the indicated message providers.

To successfully create a functional tracing configuration with one or more added system ETW Providers, you should be familiar with the workings of system ETW Providers before you employ them; however, you are free to experiment to see what results you obtain. For this reason, adding system ETW Providers to a Live Trace Session is recommended for advanced Message Analyzer users. However, Message Analyzer ships with several built-in **Trace Scenarios** that contain system ETW Providers, and these can serve as usage examples for new users. You can see a list of system ETW Providers in the **Add System Providers** dialog, which is accessible by clicking the **Add System Providers** item in the **Add Providers** drop-down list on the **ETW Providers** toolbar of the **New Session** dialog.

More Information
To learn more about how to specify one or more system ETW Providers for a Live Trace Session configuration, see Adding a System ETW Provider.
To learn more about specifying **Keyword** and **Level** filters, see System ETW Provider Event Keyword/Level Settings.

To learn more about other filtering configurations that you can apply to a built-in **Trace Scenario**, see the following topics:

Selecting Data to Capture
Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog
Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog

See Also

Configuring Session Scenarios with Selected Data Sources
Editing Existing Sessions

Targeting Live Data as an Input Source

9 minutes to read

Message Analyzer enables you to capture live traffic from network protocols, ETW-instrumented Windows system components, certain devices, and application communications. To capture this traffic, Message Analyzer makes use of several Microsoft Protocol Engineering Framework (PEF) providers that are designed to capture messages at different stack entry points (layers), which includes the capture of events at the ETW layer, along with a host of other Windows ETW-instrumented message providers that capture events from various Windows components at the ETW layer, for example, DNS, DHCP, Active Directory, and many other providers that are available from the **Add System Providers** dialog. As such, these message providers and various combinations thereof are the means of targeting live data of different types as input to Message Analyzer. Moreover, certain message providers are combined and included in built-in **Trace Scenarios** to optimize each scenario for capturing specific data. These scenarios are available for selection when you are configuring a Live Trace Session.

Using Trace Scenarios as an Input Source

The following are a few examples of built-in **Trace Scenarios** that serve as input sources for Message Analyzer. They use one or more message providers to return specific types of message data after you start a Live Trace Session:

- **Local Network Interfaces** scenario — uses the **Microsoft-PEF-NDIS-PacketCapture** provider on computers running the Windows 7, Windows 8, or Windows Server 2012 operating systems, or uses the **Microsoft-Windows-NDIS-PacketCapture** provider on computers running Windows 8.1, Windows Server 2012 R2, Windows 10, or later. Captures traffic at and above the Link Layer. Note that the **Microsoft-PEF-NDIS-PacketCapture** provider has an **Advanced Settings** dialog that enables you to configure **Fast Filter** groups. The **Microsoft-Windows-NDIS-PacketCapture** provider also has an **Advanced Settings** dialog that provides special features and capabilities, for example, specifying advanced filtering configurations when capturing traffic from remote hosts and virtual machines (VMs) that are serviced by a Hyper-V-Switch.

> **IMPORTANT**
>
> The **Microsoft-Windows-NDIS-PacketCapture** provider uses Windows Management Instrumentation (WMI) to capture remote traffic from any target host that is running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system, from any computer that is also running one of these operating systems.

- **Loopback and Unencrypted IPSEC** scenario — uses the **Microsoft-Pef-WFP-MessageProvider** to capture local traffic above the IP/Network layer, which includes transport messages such as TCP, UDP, and anything above these; this includes other messages above the IP layer, for example, those that are issued by the DNS protocol, SOAP, and RPC. Note that the **Microsoft-Pef-WFP-MessageProvider** has an **Advanced Settings** dialog that enables you to configure **Fast Filters** and **WFP Layer Set** filters.

> **IMPORTANT**
>
> The **Microsoft-Pef-WFP-MessageProvider** uses WMI to facilitate remote capabilities on Windows 10 (and later) computers only. In practice, this means that you can target a Windows 10 computer from which to capture remote traffic and return it to a computer that is running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system, where you started a **Trace Scenario** that contains the **Microsoft-Pef-WFP-MessageProvider**.

- **Pre-Encryption for HTTPS** scenario — uses the **Microsoft-Pef-WebProxy** provider, which is based on

Fiddler, to capture local HTTP client browser traffic at the Application Layer prior to HTTPS encryption. Note that the **Microsoft-Pef-WebProxy** provider has an **Advanced Settings** dialog that enables you to configure **Hostname** and **Port** filters.

> **NOTE**
>
> To enable this scenario, you must have Fiddler installed on your computer. If you do not, then go here to download and install the Fiddler Library from Telerik.

- **USB2** scenario — uses the **Microsoft-Windows-USB-USBPORT** and **Microsoft-Windows-USB-USBHUB** providers to capture events related to USB2 devices plugged into a USB port, for troubleshooting such devices.
- **SMB Client and Firewall** scenario — uses the **Microsoft-Windows-SMBClient** and **Microsoft-Pef-WFP-MessageProvider** to capture SMB2 client traffic with message headers only, along with capturing network traffic above the IP layer for correlation with the SMB2 traffic, respectively.

> **NOTE**
>
> To learn more about PEF message providers and the **Microsoft-Windows-NDIS-PacketCapture** provider, see the **More Information** section below.

Selecting a Built-In Trace Scenario

You can locate the built-in **Trace Scenarios** in the **Select Scenario** drop-down list on the **ETW Providers** toolbar of the **New Session** dialog, as shown in the figure that follows.

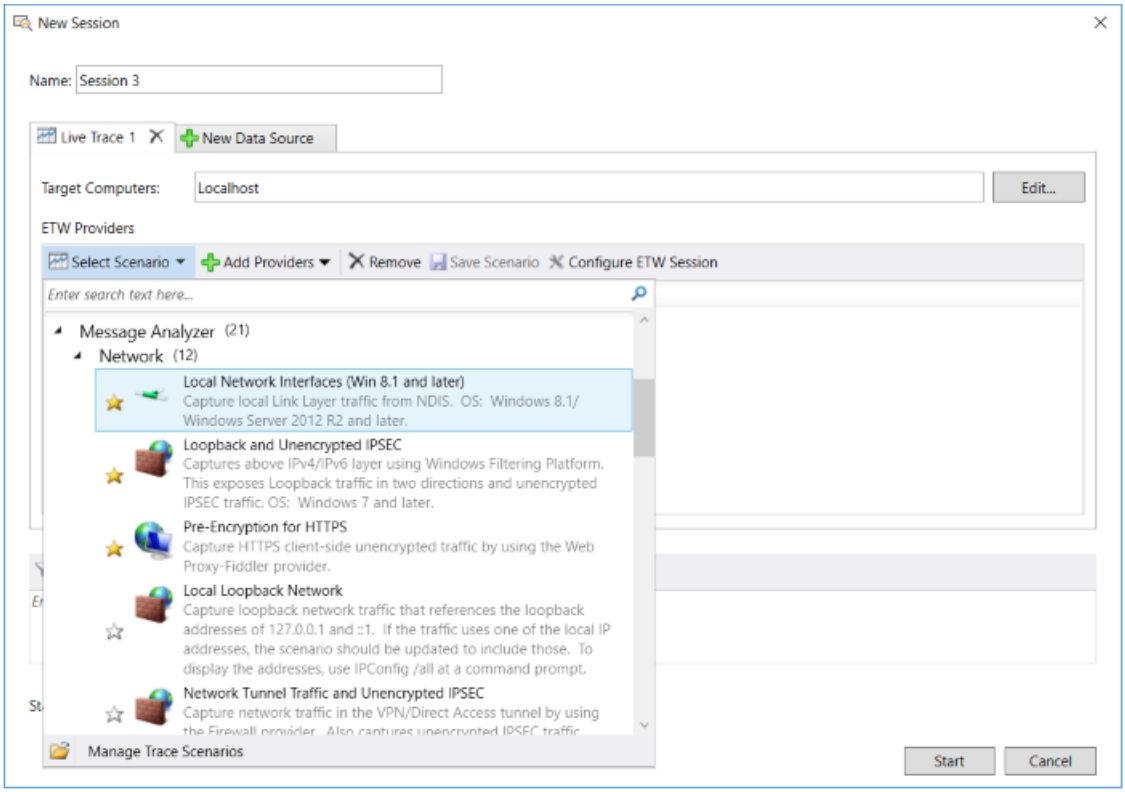

Figure 19: Message Analyzer Built-In Trace Scenarios for Live Trace Session Configuration

These **Trace Scenarios** encapsulate built-in message provider configurations that are optimized to capture specific types of data in a Live Trace Session. To add a basic **Trace Scenario** with no additional configuration required for a Live Trace Session, you can select a scenario in the **Select Scenario** drop-down list in the **New**

Session dialog. The scenario that you choose should contain the provider/s that capture data at a particular stack layer, or should focus on capturing messages from some device, application, or system component in which you are interested. Note that a description is provided with each **Trace Scenario** in the **Select Scenario** drop-down list so that you can get a sense of its functionality. After you select a scenario, you can begin capturing data immediately with a single click on the **Start** button of the **New Session** dialog. Thereafter, the message data displays in the data viewer that is set as the default, for example, the **Analysis Grid** viewer.

More Information
To learn more about the functions of the **Trace Scenarios** provided by default in every Message Analyzer installation, see the Built-In Trace Scenarios topic.

Quick Start Trace Scenarios

Message Analyzer also enables you to use a few different methods to start a live trace immediately, with no initial configuration required. These methods also make use of built-in **Trace Scenarios**, which contain message providers that are tailored to capture specific types of data, as follows:

- **Start Local Trace** — click this button on the **Start Page** to begin capturing messages at and above the Link Layer with the **Microsoft-Windows-NDIS-PacketCapture** provider in the **Local Network Interfaces** scenario on your Windows 8.1, Windows Server 2012 R2, or Windows 10 computer, or with the **Microsoft-PEF-NDIS-PacketCapture** provider in the **Local Network Interfaces** scenario on your local Windows 7, Windows 8, or Windows Server 2012 computer.
- **Favorite Scenarios** — click this drop-down list on the Message Analyzer global toolbar and then click one of the favorited **Trace Scenarios** in the list, or select a **Trace Scenario** from the **Favorite Scenarios** list on the **Start Page**. The **Trace Scenarios** that are included in the **Favorite Scenarios** list by default are the **Local Network Interfaces**, **Loopback and Unencrypted IPSEC**, and **Pre-Encryption for HTTPS** scenarios, which capture messages at and above the Link Layer, above the IP/Network Layer, and unencrypted at the Application Layer, respectively. **Favorite Scenarios** are also accessible from the Message Analyzer **File** menu. Note that you can set other **Trace Scenarios** to Favorite status as you discover their compatibility with your needs.

Using System ETW Providers as an Input Source

In a Live Trace Session, instead of capturing message data through a built-in **Trace Scenario**, you have the option to specify one or more Windows system ETW Providers from the **Add System Providers** dialog, which is accessible by clicking the **Add System Providers** item in the **Add Providers** drop-down list on the **ETW Providers** toolbar of the **New Session** dialog. Although there are no restrictions on using these providers, you should be familiar with the type of data they capture before you use them. Otherwise, it might be difficult to make sense of the data you capture.

Possible information sources that you might consult to learn more about the type of data that system ETW Providers capture consist of the following:

- **Advanced Settings** dialog — click an ellipsis (**...**) on the **ETW Core** tab of the **Advanced Settings** dialog for any message provider to view a list of **Keyword** filters that such providers make available for selection, as described in System ETW Provider Event Keyword/Level Settings.
- **Provider manifest files** — these are typically cached on your machine. The manifests typically include **Keywords** that define the events that a provider captures and can also contain event descriptions that can be useful when assessing ETW Provider functionality. To see an example of a simple manifest, see the Event Manifest topic in the ETW Framework Conceptual Tutorial.
- **Windows Events Provider Explorer (WEPExplorer)** — a graphic utility that enables you to display metadata from the manifest of each system ETW Provider that is installed and registered on your system.

The metadata includes event **Keywords**, error **Levels**, **Opcodes**, **Channels**, and so on. For further details, see Windows Events Provider Explorer on the web.

- **Event Trace Sessions** — view **Keyword** configurations for various ETW Providers in live sessions that you can view with Performance Monitor, as described in Finding System ETW Provider Metadata. If you select an ETW Provider in the **Providers** list in the **Properties** dialog for a particular session and then click the **Edit** button, you can display the list of **Keywords** that were originally specified in the manifest as the selected provider was developed. This can give you an indication of the type of messages that the selected ETW Provider will capture. In most cases, you can also locate such providers in the previously mentioned **Add System Providers** dialog when you are configuring a Live Trace Session in Message Analyzer.
- **Community knowledge bases** — you may be able to find information on optimized ETW Provider configurations from sources such as blogs and other social media.

Optimizing Data Capture Configurations

There are several ways that you can optimize your data capture configurations. Most of them involve additional filtering of some sort. Other methods include specifying a **Trace Scenario** or ETW message provider that captures specific types of messages, or messages at different Layers, as previously indicated. The goal is to maintain optimal performance while capturing the least amount of data to solve the problem at hand. For example, by specifying a **Session Filter**, ETW **Keyword** filter, **Fast Filter**, **Parsing Level**, a **Trace Scenario** that focuses on messages at a particular Layer, or by applying other advanced filtering techniques, you can optimize the process of capturing data in a Live Trace Session to achieve the following:

- Reduced message count in a more manageable data set that focuses on specific data only.
- Faster performance and reduced use of system resources when capturing message data.
- A more streamlined data analysis process.
- A focused set of messages that expedites the analysis process for others with whom you are sharing your results.

More Information
To learn more about using **Session Filters**, see Working with Session Filters in a Live Trace Session.
To learn more about using **Keyword** filters, see Adding a System ETW Provider and System ETW Provider Event Keyword/Level Settings.
To learn more about using **Fast Filters**, see PEF-NDIS Fast Filters and PEF-WFP Fast Filters.
To learn more about using **Parsing Levels**, see Setting the Session Parsing Level.
To learn more about the capabilities of the **Microsoft-Windows-NDIS-PacketCapture** provider, see the Microsoft-Windows-NDIS-PacketCapture Provider topic.

Session Configuration Workflow and Features

The linked section immediately below provides a general workflow that you can follow when configuring a Live Trace Session. It also contains subtopics that describe the features and functions that you can use when configuring a Live Trace Session:
Configuring a Live Trace Session

See Also

Built-In Trace Scenarios
PEF Message Providers

PEF Message Providers

2 minutes to read

Message Analyzer can make use of many different message providers when capturing data. Certain providers are specifically designed by Microsoft to create inspection points into the network protocol stack so that you can focus on retrieving and analyzing messages at predefined levels. For example, you can use the **Microsoft-PEF-WFP-MessageProvider** to retrieve network traffic above the IP/Network layer while minimizing lower level network noise, so that you can focus on troubleshooting TCP and application messages. You can also use the **Microsoft-PEF-NDIS-PacketCapture** provider to retrieve network traffic on the wire at the Data Link layer to obtain a full representation of all messages captured in a trace. In addition, you can use the **Microsoft-PEF-WebProxy** provider to capture HTTP client browser traffic at the Application Layer, prior to encryption.

This section briefly describes the functions of Microsoft Protocol Engineering Framework (PEF) providers that are native to Message Analyzer and which enable you to partition the capture of network traffic in the indicated manner. The filters that you can configure for these providers are also described in the sections below:

Microsoft-PEF-NDIS-PacketCapture Provider
Microsoft-PEF-WFP-MessageProvider
Microsoft-PEF-WebProxy Provider

TIP

Note that you can also use the **Microsoft-Windows-NDIS-PacketCapture** provider to capture messages locally or remotely at Link Layer.

About Provider Manifests

Each Microsoft PEF provider has an ETW manifest that installs with Message Analyzer. A provider manifest is an XML file that specifies a formal description of the events a provider raises. It identifies the event provider, specifies the event types, and also describes the events.

A manifest can also associate its events with **Keywords** and **Levels**, which is a way to enable events and filter them as they are written for consumption:

- **Keywords** — group events together that are logically related.
- **Level** — indicates the severity or verbosity of an event, for example, critical, error, warning, or informational.

TIP

Keywords are different for many ETW providers. You might therefore consider consulting the community knowledge base for optimized configurations.

In addition, event consumers such as the PEF Runtime can make use of a manifest's structured XML data to perform queries and analysis. Manifests for all PEF providers reside in the following location:

```
c:\Windows\System32\
```

See Also

Event Manifest

Microsoft-PEF-NDIS-PacketCapture Provider

4 minutes to read

The **Microsoft-PEF-NDIS-PacketCapture** filter driver is primarily used by the Message Analyzer **Local Network Interfaces Trace Scenario** on computers running the Windows 7, Windows 8, and Windows Server 2012 operating systems. It is instrumented to work with the ETW infrastructure, which provides the mechanisms for controlling ETW Sessions, buffering data, and delivering captured events to a consumer. Because the **Microsoft-PEF-NDIS-PacketCapture** filter driver works with this infrastructure, it can deliver the frames it captures as ETW events. In addition to enabling the capture of frames, the **Microsoft-PEF-NDIS-PacketCapture** provider also inspects and infuses frames at the Data Link miniport level. Each network interface is represented by a miniport driver that performs Data Link Layer frame operations.

Provider Functions

The **Microsoft-PEF-NDIS-PacketCapture**[1] provider is a light-weight filter (LWF) Network Driver Interface Specification (NDIS) driver that uses the Data Link Layer as its entry point into the message stack. It captures frames at the Data Link Layer and stack messages above that level, in addition to low level events from the underlying ETW provider with which the **Microsoft-PEF-NDIS-PacketCapture** driver is instrumented. Because it is instrumented for ETW, the **Microsoft-PEF-NDIS-PacketCapture** provider can take advantage of the ETW infrastructure to deliver its events and messages to an enabling ETW Session. The ETW Session model consists of interaction between the following three components:

- ETW Event Provider
- ETW Session Controller
- ETW Event Consumer

The ETW event provider in this model is the **Microsoft-PEF-NDIS-PacketCapture** filter driver, while the ETW event consumer is the Protocol Engineering Framework (PEF) Runtime, which parses/processes the message and event data that is delivered by the **Microsoft-PEF-NDIS- PacketCapture** provider and exposes the processed data in an API that the Message Analyzer user interface (UI) consumes for the display of data.

Associated Trace Scenarios

The **Microsoft-PEF-NDIS-PacketCapture** provider is used in the following Message Analyzer **Trace Scenarios**, where each scenario has a different capture focus. The **Microsoft-PEF-NDIS-PacketCapture** provider is available in these scenarios on computers running the Windows 7, Windows 8, or Windows 2012 operating system only:

- **Local Network Interfaces**
- **VPN**
- **Wired Local Area Network**
- **Wireless Local Area Network**

Provider Fast Filtering Configurations

The **Microsoft-PEF-NDIS-PacketCapture** driver also provides a filtering mechanism known as **Fast Filters**. These filters enable you to apply frame filtering based on offset length patterns (OLPs), MAC addresses

(LinkLevelAddress), and IP addresses. Because **Fast Filters** operate at the driver level, they are typically very efficient and performant. However, a disadvantage of these filters is that you cannot recover any data that was filtered out of a Live Trace Session by a **Fast Filter**. **Fast Filters** for the **Microsoft-PEF-NDIS-PacketCapture** provider are described in PEF-NDIS Fast Filters.

> **TIP**
> One exception to recovering data that was previously filtered-out is when you have a Message Analyzer Data Retrieval Session in which a **Session Filter** was applied. In this case, after you load data into Message Analyzer with a **Session Filter** applied, you can reopen the session configuration and remove the existing **Session Filter** and then reload the data with no filtering applied, as described in Editing Existing Sessions.

You can specify up to three **Fast Filters** in two logically chained Groups for any Live Trace Session that uses the **Microsoft-PEF-NDIS-PacketCapture** provider. Note that you can use up to four **Fast Filters** for Live Trace Sessions that use the **Microsoft-PEF-WFP-MessageProvider**. Frames that pass the **Fast Filter** criteria are delivered to the enabling ETW Session as events.

When configuring a Live Trace Session that uses the **Microsoft-PEF-NDIS-PacketCapture** provider, you can specify **Fast Filter** settings on the **Provider** tab of the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, which is accessible by clicking the **Configure** link next to the provider **Id** (GUID) in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog. You can access this dialog from the Message Analyzer **File** menu, the global Message Analyzer toolbar, or from the **Start Page**.

Event Keyword and Error Level Filtering

Similar to other system ETW providers that are registered on your system, you have the option to specify ETW event **Keyword** configurations and **Level** filter settings for the **Microsoft-PEF-NDIS-PacketCapture** provider, to further refine the scope of events to be captured. The settings for these filters are available on the **ETW Core** tab of the **Advanced Settings – Microsoft-PEF-NDIS-PacketCapture** dialog. By default, all events are delivered by the **Microsoft-PEF-NDIS-PacketCapture** provider when no **Keywords** are selected. If you specify any particular **Keyword**, then the provider will deliver the events that are enabled by that **Keyword** only, if they occur in a trace. It also follows that multiple events are delivered when multiple **Keywords** are selected.

More Information
To learn more about configuring system ETW Providers, including event **Keyword** and error **Level** filters, see System ETW Provider Event Keyword/Level Settings.
To learn more about configuring PEF providers, including setting **Fast Filters** for the **Microsoft-PEF-NDIS-PacketCapture** provider, see the following topics:
Common Provider Configuration Settings Summary
Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog
To learn more about the ETW session model, see the ETW Framework Conceptual Tutorial.

---

[1] Computers that are running the Microsoft Windows 8 and Microsoft Windows Server 2012 64-bit operating systems use the Microsoft-PEF-NDIS-PacketCapture v6.3 provider. Computers running the Microsoft Windows 7 operating system or the 32-bit version of the Microsoft Windows 8 operating system use the Microsoft-PEF-NDIS-PacketCapture v6.0 provider. Computers running the Microsoft Windows 8.1, Windows Server 2012 R2, Windows 10, and later operating systems use the Microsoft-Windows-NDIS-PacketCapture (NDISCAP) provider only, which has remote capabilities.

See Also

Built-In Trace Scenarios
Microsoft-PEF-WFP-MessageProvider

Microsoft-PEF-WebProxy Provider

PEF-NDIS Fast Filters

2 minutes to read

The **Microsoft-PEF-NDIS-PacketCapture** filter driver enables you to select messages from a trace that meet certain criteria. You can select specific message data from a trace by configuring a **Fast Filter** on the **Provider** tab of the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, which is accessible by clicking the **Configure** link to the right of the provider **Id** in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog. Messages that pass the filtering criteria are then passed to Message Analyzer Runtime processing. Driver-level filtering can lower system loads by passing less data, reducing CPU processing, and writing fewer events, which prevents costly disk I/O. Driver filtering enables significant improvements in speed over copying messages to the Runtime and then filtering with its parsing engine.

> **TIP**
>
> Following the use of a **Fast Filter** in a Live Trace Session, you still have the option to further filter your trace results with a view **Filter**, for analysis purposes. You can specify a view **Filter** from the Filter panel **Library** that displays when you click **Add Filter** in the **Add Filter** drop-down list on the Filtering toolbar.

The different types of **Fast Filters** for the **Microsoft-PEF-NDIS-PacketCapture** provider are described in the following topics:

OLP Filters
LinkLevelAddress Filters
IPv4Address Filters
IPv6Address Filters

By using these filters, you can direct the **Microsoft-PEF-NDIS-PacketCapture** provider to isolate message traffic based on values of specific types of addresses or offset length patterns (OLPs).

More Information

To learn more about the settings for **Microsoft-PEF-NDIS-PacketCapture Fast Filters**, see the Common Provider Configuration Settings Summary.

To learn more about **Fast Filter** configuration capabilities, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

OLP Filters

2 minutes to read

An OLP filter is a type of **Fast Filter** that enables you to locate message fields containing a specific **IPv4Address**, **IPv6Address**, **LinkLevelAddress**, string, or other value, by specifying an **O**ffset value, bit **L**ength, and value **P**attern. Only messages containing a match to the specified pattern are identified by the OLP filter for further processing. Messages containing a match are then either included or excluded from the returned trace data, depending on the operator you configure in the OLP expression.

Using the OLP Format for Fast Filters

You can use the following format as a template when specifying an OLP filter as a **Fast Filter**:
op bit-offset:bit-length:hex-value
where "op" can be one of the following operators:

- Equals (==). If not specified, "==" is assumed.
- Not equal to (!=).
- Greater than (>).
- Less than (<).

The OLP filter parameters for which you specify values are delimited by a colon (:). These parameters have the following meanings:

- **Bit-offset** — is specified in hexadecimal format and represents the number of offset bytes that precede an address being searched for.
- **Bit-length** — is specified as an integer that equals the length in bits of the hex-value pattern.
- **Hex-value** — is specified in hexadecimal format and represents the address pattern for which you are searching.
  Address patterns can be any of the following:
  - An **IPv4Address** in x.x.x.x format — is specified as an 8-digit hexadecimal number, 32-bits long; for example: *9D3B54DF*.
  - An **IPv6Address** in xx:xx:xx:xx:xx:xx:xx:xx format — is specified as a 32-digit hexadecimal number, 128-bits long; for example: *9D3B54DFC4D7A46F0000000000000000*.
  - A **LinkLevelAddress** in xx-xx-xx-xx-xx-xx format — is specified as a 12-digit hexadecimal number, 48-bits long; for example: *001F297D2F46*.

By using an OLP filter, you can be very specific and granular about pinpointing messages that meet the filtering criteria that you define. This can have a big impact on reducing the amount of traffic you capture and improving performance, for example, to minimize the effect of a capture on a busy server.

Configuring an OLP Fast Filter

To configure an **OLP** filter, you must open the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog by clicking the **Configure** link to the right of the provider **Id** in the **ETW Providers** list that displays on the **Live Trace** tab of the **New Session** dialog. When the dialog displays, select the **Provider** tab and then click a **Filter** drop-down arrow in a **Fast Filter Group**. In the menu that appears, select the **OLP** item as the address type.

You must then enter an OLP address in the text box to the right of the selected filter type, in the format that is described in this section.

You have the option to configure up to three **Fast Filters** per **Group** in the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, with the same or different address types. When you are finished with **Fast Filter** configuration, highlight the **System Network** tree grid row containing the adapter that you want to assign the **OLP Fast Filter** to, click the **Apply To Highlighted** button to assign the filter **Group** to the adapter, and then click **OK** to close the dialog. You can then start your Live Trace Session, at which time the **Microsoft-PEF-NDIS-PacketCapture** provider automatically isolates and captures messages that meet the criteria of the filtering configuration that you specified.

More Information

To learn more about **Fast Filter** configuration capabilities, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

LinkLevelAddress Filters

2 minutes to read

A **LinkLevelAddress** is a type of **Fast Filter** that enables you to filter out all messages during a Live Trace Session except those that are sent to and from a specified physical address. You can also use several relational operators with **Fast Filters** to enhance filtering functionality. For example, for **LinkLevelAddress** filters, you could use the logical NOT (!=) operator to filter out all messages that are sent to and from a specified physical address. You might do this if you have multiple network adapters and you want to isolate traffic to a particular one. Other operators that are available for use with **Fast Filters** include EQUALS (==), LESS THAN (<), and GREATER THAN (>).

Configuring a LinkLevelAddress Fast Filter

To configure a **LinkLevelAddress** filter, you must open the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog by clicking the **Configure** link to the right of the **Microsoft-PEF-NDIS-PacketCapture** provider **Id** in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog. When the **Advanced Settings** dialog displays, select the **Provider** tab and then click a **Filter** drop-down arrow in a **Fast Filter Group**. In the menu that appears, select the **LinkLevelAddress** item as the address type. You must then enter a media access control (MAC) address for the network adapter on which to capture messages by specifying its value in the text box to the right of the selected filter type, in a format similar to the following examples:

```
00-2F-39-7E-1F-36
!=00-2F-39-7E-1F-36
```

You have the option to configure up to three **Fast Filters** per **Group** in the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, with the same or different address types. For example, you might want to use three MAC addresses to collect or block data from a specified set of network adapters. When you are finished with **Fast Filter** configuration, highlight the **System Network** tree grid row containing the adapter that you want to assign the **LinkLevelAddress Fast Filter** to, click the **Apply To Highlighted** button to assign the filter **Group** to the adapter, and then click **OK** to close the dialog. You can then start your Live Trace Session, at which time the **Microsoft-PEF-NDIS-PacketCapture** provider automatically isolates and captures messages that meet the criteria of the filtering configuration that you specified.

More Information

To learn more about **Fast Filter** configuration capabilities, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

IPv4Address Filters

2 minutes to read

An **IPv4Address** is a type of **Fast Filter** that enables you to filter out all messages during a live trace except those that are sent to or from a specified IPv4 address. As with other **Fast Filter** configurations, you can also use several relational operators to enhance filtering functionality. For example, for **IPv4Address** filters, you could use the logical NOT (!=) operator to filter out all messages that are sent to and from a specified IPv4 address. You could also use the GREATER THAN (>) or LESS THAN (<) operators to isolate message traffic above or below a particular IPv4 address value, respectively. You might do this if you want to isolate traffic to a particular subnet.

Configuring an IPv4Address Fast Filter

To configure an **IPv4Address** filter, you must open the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog by clicking the **Configure** link to the right of the **Microsoft-PEF-NDIS-PacketCapture** provider **Id** in the **ETW Providers** list that displays on the **Live Trace** tab of the **New Session** dialog. When the **Advanced Settings** dialog displays, select the **Provider** tab and then click a **Filter** drop-down arrow in a **Fast Filter Group**. In the menu that appears, select the **IPv4Address** item as the address type. You must then enter the IPv4 address value with or without operators, by specifying it in the text box to the right of the selected filter type, in a format similar to the following examples:

```
192.168.1.1
```

```
!= 192.168.1.1
```

```
< 192.168.1.1
```

You have the option to configure up to three **Fast Filters** per **Group** in the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, with the same or different address types. For example, you might want to collect or block data from several IPv4 addresses. When you are finished with **Fast Filter** configuration, highlight the **System Network** tree grid row containing the adapter that you want to assign the **IPv4Address Fast Filter** to, click the **Apply To Highlighted** button to assign the filter **Group** to the adapter, and then click **OK** to close the dialog. You can then start your Live Trace Session, at which time the **Microsoft-PEF-NDIS-PacketCapture** provider automatically isolates and captures messages that meet the criteria of the filtering configuration that you specified.

More Information

To learn more about **Fast Filter** configuration capabilities, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

IPv6Address Filters

2 minutes to read

An **IPv6Address** is a type of **Fast Filter** that enables you to filter out all messages during a live trace except those that are sent to and from a specified IPv6 address. As with other **Fast Filter** configurations, you can also use several relational operators to enhance filtering functionality. For example, you could use the logical NOT (!=) operator to filter out all messages that are sent to and from a specified IPv6 address. You can also use the GREATER THAN (>) or LESS THAN (<) operators to isolate message traffic above or below a particular IPv6 address value, respectively.

Configuring an IPv6Address Fast Filter

To configure an **IPv6Address** filter, you must open the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog by clicking the **Configure** link to the right of the **Microsoft-PEF-NDIS-PacketCapture** provider **Id** in the **ETW Providers** list that displays on the **Live Trace** tab of the **New Session** dialog. When the **Advanced Settings** dialog displays, select the **Provider** tab and then click a **Filter** drop-down arrow in a **Fast Filter Group**. In the menu that appears, select the **IPv6Address** item as the address type. You must then enter the IPv6 address value with or without operators, by specifying it in the text box to the right of the selected filter type, in a format similar to the following examples:

```
2001:4898:2b:3:ac0a:511:f971:5d85
```

```
!= 2001:4898:2b:3:ac0a:511:f971:5d85
```

```
> 2001:4898:2b:3:ac0a:511:f971:5d85
```

You have the option to configure up to three **Fast Filters** per **Group** in the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, with the same or different address types. For example, you might want to collect or block data from several IPv6 addresses. When you are finished with **Fast Filter** configuration, highlight the **System Network** tree grid row containing the adapter that you want to assign the **IPv6Address Fast Filter** to, click the **Apply To Highlighted** button to assign the filter **Group** to the adapter, and then click **OK** to close the dialog. You can then start your Live Trace Session, at which time the **Microsoft-PEF-NDIS-PacketCapture** provider automatically isolates and captures messages that meet the criteria of the filtering configuration that you specified.

More Information

To learn more about **Fast Filter** configuration capabilities, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

PEF-NDIS Provider Manifest

2 minutes to read

The manifest for the **Microsoft-PEF-NDIS-PacketCapture** provider defines its contract with an ETW Controller in terms of the events that are issued by the provider, which are in turn consumed by Message Analyzer. The **Microsoft-PEF-NDIS-PacketCapture** provider and its manifest are installed and registered on your machine when you install Message Analyzer. You can find the **Microsoft-PEF-NDIS-PacketCapture** provider manifest file in the following location:

```
C:\Windows\System32\Microsoft-Pef-NDIS-PacketCapture.man
```

In this file you can view the event definitions, tasks, opcodes, keyword bitmask and level specifiers, a channel definition, and template information.

> **NOTE**
> You will find the Microsoft-Pef-NDIS-PacketCapture.man file in the specified location only after installing Message Analyzer on computers that are running the Windows 7, Windows 8, or Windows Server 2012 operating system.

More Information

To learn more about the contents of provider manifests, see Event Manifest.

To learn more about how Message Analyzer uses provider manifests, see Understanding Event Parsing with a Provider Manifest.

To learn more about obtaining a provider manifest if you are missing one, see Generating a Provider Manifest.

Microsoft-PEF-WFP-MessageProvider

5 minutes to read

Similar to the **Microsoft-PEF-NDIS-PacketCapture** provider, the **Microsoft-PEF-WFP-MessageProvider** is instrumented to work with the ETW infrastructure, as described in the ETW Framework Conceptual Tutorial. This infrastructure provides the mechanisms for controlling ETW Sessions, buffering trace data, and delivering events to a consumer. Because the **Microsoft-PEF-WFP-MessageProvider** provider is integrated into this infrastructure, it can deliver the packets that it captures as events. The **Microsoft-PEF-WFP-MessageProvider** also has an option that enables you to log discarded packet events.

Provider Functions

The **Microsoft-PEF-WFP-MessageProvider** is based on the Windows Filtering Platform (WFP) which is typically used to create firewall applications, but in Message Analyzer, it provides an entry point into the stack that enables you to capture traffic above the IP/Network Layer. Capturing message data with the **Microsoft-PEF-WFP-MessageProvider** is less noisy than capturing at the Data Link Layer with either the **Microsoft-PEF-NDIS-PacketCapture** or **Microsoft-Windows-NDIS-PacketCapture** provider, as most low-level and broadcast traffic is ignored in the case of the **PEF-WFP-MessageProvider**. Therefore, with the use of the **Microsoft-PEF-WFP-MessageProvider** you can focus your analysis on messages above the Network Layer. Note that messages at and below the Network Layer in a **Microsoft-PEF-WFP-MessageProvider** capture are typically represented in your trace results as **WFPCapture**, and below that there is an **ETW** layer.

Although most lower-level noise is removed by the **Microsoft-PEF-WFP-MessageProvider**, other noise can be introduced by this provider in terms of diagnosis errors and loopback traffic. However, you can specifically filter out this noise by configuring **Fast Filters** in the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog, as described in Provider Fast Filtering Configurations. The provider configuration in this dialog is accessible by clicking the **Configure** link in the **New Session** dialog in **Trace Scenarios** that use this provider.

The **Microsoft-PEF-WFP-MessageProvider** provides the first entry point (chokepoint) into the network stack above the Data Link Layer in Message Analyzer, which creates a unique inspection point into the stack that enables you to do the following:

- **Capture loopback traffic** — for example, if you have an HTTP and SQL server on the same machine, you can use the **Local Loopback Network Trace Scenario** to capture communication traffic between these applications, that is, if they are using the loopback IP addresses for such communications.
- **Capture IPSec/ESP traffic** — in pre-encrypted, encrypted, and decrypted formats, for example, with the **Loopback and Unencrypted IPSEC Trace Scenario**.
- **Capture VPN and Direct Access traffic** — for example, with the **Network Tunnel Traffic and Unencrypted IPSEC Trace Scenario**.
- **Capture hardware offload traffic** — offloaded from the network adapter.
- **Log discarded packet events** — by setting the **Select Discarded Packet Events** option, Message Analyzer can capture packets that are being dropped at the firewall.
- **Improve performance** — use less Message Analyzer parsing to get at upper-layer protocols.

Associated Trace Scenarios

The **Microsoft-PEF-WFP-MessageProvider** provider is used in the following Message Analyzer **Trace**

Scenarios, where each scenario has a different capture focus. Operating system dependencies are also specified where applicable; otherwise, you can assume that the **Microsoft-PEF-WFP-MessageProvider** works in the listed **Trace Scenarios** on Windows 7 and later computers:

- **Local Loopback Network**
- **Loopback and Unencrypted IPSEC**
- **Network Tunnel Traffic and Unencrypted IPSEC**
- **SASL LDAP pre-encryption with WFP**
- **SMB Client and Firewall** — on computers running the Windows 8, Windows Server 2012, or Windows 10 operating system.

> **IMPORTANT**
>
> Note that you can use any **Trace Scenario** in which the **Microsoft-PEF-WFP-MessageProvider** exists on computers running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system, to capture traffic remotely on computers running the Windows 10 operating system. Capture of remote traffic on target Windows 10 computers is possible because the **Microsoft-PEF-WFP-MessageProvider** is now instrumented to support remote scenarios on Windows 10 computers.
>
> This means that on computers running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system, you can successfully launch a remote capture that targets a Windows 10 computer, while on computers with other operating systems earlier than Windows 8.1, you cannot. Conversely, from a computer running the Windows 10 operating system, you cannot capture remote traffic from a computer running the Windows 8.1, Windows Server 2012 R2, or earlier operating system, since the **Microsoft-PEF-WFP-MessageProvider** is not instrumented to support remote scenarios on these computers. Moreover, it is the target computer with the remote-enabled **Microsoft-PEF-WFP-MessageProvider** alone that determines where you can capture remote traffic.

Provider Fast Filtering Configurations

Similar to the **Microsoft-PEF-NDIS-PacketCapture** provider, you can configure the **Microsoft-PEF-WFP-MessageProvider** to use **Fast Filters** to improve trace performance. For example, you might create a **Fast Filter** that focuses on messages from a specified IP address or TCP port only. You can also specify **Fast Filters** that remove loopback traffic, for example, with **IPv4** and **IPv6** filters such as `!127.0.0.1` and `!::1`, respectively. You can access the configuration for **Fast Filters** on the **Provider** tab of the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog, which is accessible by clicking the **Configure** link to the right of the **Microsoft-PEF-WFP-MessageProvider Id** in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog.

> **NOTE**
>
> The **WFP Layer Set** filter configuration is also specified on the **Provider** tab of the **Advanced Settings** dialog, as described in the PEF-WFP Layer Set Filters topic.

Event Keyword and Error Level Filtering

Similar to other system ETW providers that are registered on your system, you have the option to specify ETW event **Keyword** configurations and **Level** filter settings for the **Microsoft-PEF-WFP-MessageProvider**, to further refine the scope of events to be captured. The settings for these filters are available on the **ETW Core** tab of the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog. By default, as with the **Microsoft-PEF-NDIS-PacketCapture** provider, all events are delivered by the **Microsoft-PEF-WFP-MessageProvider** provider when no **Keywords** are selected. If you specify any particular **Keyword**, then the provider will deliver the events that are enabled by that **Keyword** only, if they occur in a trace. It also follows that multiple events are

delivered when multiple **Keywords** are selected.

More Information

To learn more about configuring system ETW Providers, including event **Keyword** and error **Level** filters, see System ETW Provider Event Keyword/Level Settings

To learn more about specifying **Fast Filters** and **WFP Layer Set** filters for the **Microsoft-PEF-WFP-MessageProvider**, see the following topics:

PEF-WFP Fast Filters

PEF-WFP Layer Set Filters

Common Provider Configuration Settings Summary

--

See Also

Built-In Trace Scenarios

Microsoft-PEF-NDIS-PacketCapture Provider

Microsoft-PEF-WebProxy Provider

PEF-WFP Layer Set Filters

2 minutes to read

When configuring settings for the **Microsoft-PEF-WFP-MessageProvider** in **Trace Scenarios** that use it, you can create a **WFP Layer Set** filter configuration that enables you to directionally isolate inbound or outbound TCP packets at the Transport layer for IPv4 or IPv6 traffic. You can access the configuration for the **WFP Layer Set** filters on the **Provider** tab of the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog, which is accessible by clicking the **Configure** link to the right of the **Id** for the **Microsoft-PEF-WFP-MessageProvider** in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog, which is accessible from the Message Analyzer **File** menu, the **Start Page**, or the global Message Analyzer toolbar. The **WFP Layer Set** contains the following filters, which you can enable or disable, respectively, by selecting or unselecting filter check boxes as appropriate:

- INBOUND_TRANSPORT_V4
- OUTBOUND_TRANSPORT_V4
- INBOUND_TRANSPORT_V6
- OUTBOUND_TRANSPORT_V6

These filters are kernel mode TCP/IP stack filters that operate in the receive or send path (inbound or outbound, respectively) at the Transport Layer before any processing occurs at that layer. When set, these filters selectively enable or disable the capture of all inbound, outbound, or bidirectional packet traffic at the Transport Layer.

Capturing Loopback Traffic

If you are capturing loopback (local application) traffic, you should disable either inbound or outbound traffic with **WFP Layer Set** filters in the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog, as the default configuration of the **Local Loopback Network Trace Scenario** does; otherwise, you will get duplicate messages. However, for regular network traffic, you should always enable both inbound and outbound packet directions.

More Information
To learn more about configuring **WFP Layer Set** filters for the **Microsoft-PEF-WFP-MessageProvider**, see the Common Provider Configuration Settings Summary.

PEF-WFP Dropped Packets

2 minutes to read

The **Microsoft-PEF-WFP-MessageProvider** enables you to log dropped packet information, which includes reason and layer statistics. To obtain these statistics, you must select the **Select Discarded Packet Events** check box on the **Provider** tab of the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog. You can access this dialog by clicking the **Configure** link to the right of the **Microsoft-PEF-WFP-MessageProvider** in the **ETW Provider** list on the **Live Trace** tab of the **New Session** dialog, that is, after selecting any **Trace Scenario** that uses the **Microsoft-PEF-WFP-MessageProvider**, for example, the **Loopback and Unencrypted IPSEC Trace Scenario**.

Logged information for discarded packet events consists of dropped packet statistics that are derived from the Discard filter layer of the Statistics callout in the **Microsoft-PEF-WFP-MessageProvider**. Statistics consist of the reason for dropping packets and the layer on which they were dropped. You can view the dropped packet statistics as ETW events in the **Analysis Grid** viewer after a trace with the **Microsoft-PEF-WFP-MessageProvider** is complete. You can use the **Column Filter** feature for the **Summary** column of the **Analysis Grid** viewer to specify a search term such as "discard" to expose any messages that might indicate that the firewall was involved in blocking traffic. The **Discarded Packet Events** feature can help you troubleshoot whether packets are being dropped by the network, the **Microsoft-PEF-WFP-MessageProvider**, or the firewall.

> **NOTE**
>
> When you select the **Select Discarded Packet Events** check box on the **Provider** tab of the **Advanced Settings** dialog, any **Fast Filters** or **WFP Layer Set** filters that you have specified will not be applied to the discarded packet events.

More Information

To learn more about using **Column Filters**, see Filtering Column Data.

PEF-WFP Fast Filters

3 minutes to read

The **Microsoft-PEF-WFP-MessageProvider** configuration on the **Provider** tab of the **Advanced Settings - Microsoft-PEF-WFP-MessageProvider** dialog contains **Fast Filter** settings, which enable you to apply filtering to messages before they are passed to the Message Analyzer Runtime for parsing. By applying a **Fast Filter** at the driver-level rather than using a **Session Filter** that is applied after parsing, you can lower system loads by reducing CPU processing, parsing less data, and writing fewer events to prevent more costly disk I/O. Driver-level filtering enables significant improvements in speed as compared to a **Session Filter**, where the Runtime parsing engine does the filtering.

Using WFP Fast Filters

The actual work that is performed by the **Fast Filters** that you specify in the **Microsoft-PEF-WFP-MessageProvider** configuration is accomplished by the WFP base filtering engine (BFE). The message frames that pass the filtering criteria are delivered to the **Microsoft-PEF-WFP-MessageProvider** callout drivers at the corresponding layers, which in turn send the messages to the enabling ETW session.

WFP Fast Filters consist of the following types:

- **IPv4** — enables you to filter live traffic based on a specified IPv4 address, as described in IPv4Address Filters.
- **IPv6** — enables you to filter live traffic based on a specified IPv6 address, as described in IPv6Address Filters.
- **TCP port** — enables you to filter live traffic based on a specified TCP port.
- **UDP port** — enables you to filter live traffic based on a specified UDP port.

When configuring **Fast Filters** in the **Advanced Settings - Microsoft-PEF-WFP-MessageProvider** dialog, you have the option to enhance filter functionality with the use of several relational operators, including logical NOT (!), GREATER THAN (>), and LESS THAN (<). For example, you might apply **Fast Filter** configurations that exclude or include specific traffic on a **TCP port** or **UDP port**, respectively. To do this, you could configure one **Fast Filter** with the **TCP port** set to remove HTTPS traffic from your trace results and another **Fast Filter** with the **UDP port** set to exclude all traffic except LDAP messages, respectively, as follows:

Fast Filter 1 value for **TCP port**: `!= 443`
Fast Filter 2 value for **UDP port**: `== 389`

Capturing Loopback Traffic

If you want to see only loopback (local application) traffic for IPv4 and IPv6 addresses in a Live Trace Session, you will need to explicitly filter for such traffic, or you can use the **Local Loopback Network Trace Scenario**, which does this with the following preset **Fast Filters**: `127.0.0.1` for the **IPv4** filter and `::1` for the **IPv6** filter.

Removing Loopback Traffic

If you want to remove local loopback traffic, you will need to explicitly filter out the loopback address and any traffic that is either coming from or going to an IPv4 or IPv6 address that is being used by the related application communications. To do this, you must negate the loopback traffic. When you use negation in a **Fast Filter** expression, it behaves similar to the way the tilde (~) character does in a **Session Filter** or view **Filter**, in that non-

existence of a specified value is not evaluated as negation. For example, the IPv4 **Fast Filter**: `!=192.168.1.1` is equivalent to the view **Filter**: `*Address ~=192.168.1.1`. With either of these filters, only IPv4 traffic is evaluated and all other traffic is ignored, whereas the view **Filter**: `*Address !=192.168.1.1` filters out the specified IP address, but also evaluates all IPv4 traffic and passes all addresses that are not 192.168.1.1.

Therefore, to filter out all loopback traffic, use **Fast Filters** similar to the following:

`!=127.0.0.1` for the **IPv4** filter and `!=::1` for the **IPv6** filter.

If the related applications are using other IP addresses, filter those out as well with more **Fast Filters**. You can specify up to four filters in the **Advanced Settings** dialog for the **Microsoft-PEF-WFP-MessageProvider**.

More Information

To learn more about configuring **Fast Filters** for the **Microsoft-PEF-WFP-MessageProvider**, see Using the Advanced Settings- Microsoft-PEF-WFP-MessageProvider Dialog.

To learn more about the Windows Filtering Platform (WFP), see Windows Filtering Platform.

PEF-WFP Provider Manifest

2 minutes to read

The manifest for the **Microsoft-PEF-WFP-MessageProvider** defines its contract with an ETW Controller in terms of the events that are issued by the provider, which are in turn consumed by Message Analyzer. The **Microsoft-PEF-WFP-MessageProvider** and its manifest are installed and registered on your machine when you install Message Analyzer. You can find the **Microsoft-PEF-WFP-MessageProvider** manifest in the following location:

```
C:\Windows\System32\Microsoft-Pef-WFP-MessageCapture.man
```

In this file you can view the event definitions, tasks, opcodes, keyword bitmask and level specifiers, a channel definition, and template information.

More Information

To learn more about the contents of provider manifests, see Event Manifest.

To learn more about how Message Analyzer uses provider manifests, see Understanding Event Parsing with a Provider Manifest.

To learn more about obtaining a provider manifest if you are missing one, see Generating a Provider Manifest.

Microsoft-PEF-WebProxy Provider

2 minutes to read

The **Microsoft-PEF-WebProxy** provider enables you to focus on capturing traffic at the Application layer of the network stack. It is based on Fiddler and enables you to capture unencrypted HTTP traffic to and from a client web browser on computers running the Windows 7 and later operating systems. However, the **PEF-WebProxy** provider will not capture unencrypted HTTP browser traffic unless you configure **Internet options** to use a proxy server for the LAN.

> **IMPORTANT**
>
> To use the **PEF-WebProxy** provider, you must have the Fiddler library from Telerik installed. If you have not already installed this library, you can download it here. Note that if you select the **WebProxy** provider without having this library installed when you are configuring a Live Trace Session, or when specifying it in a PowerShell script, you will receive an error message that asks you to download the Fiddler library.

Provider Functions

The **Microsoft-PEF-WebProxy** provider is used by the Message Analyzer **Pre-Encryption for HTTPS Trace Scenario** to capture client browser traffic. When you use this provider, it also installs and registers the HTTP proxy on the machine that is running Message Analyzer, where it is configured as a system proxy for Microsoft Windows Internet (WinInet) Services. The HTTP proxy is then substituted for the proxy server when a trace is started so that it can intercept and capture unencrypted HTTPS traffic. The proxy server target is restored when the trace is complete.

As the system proxy, all HTTPS requests from WinInet flow through the HTTP proxy before reaching target web servers. Similarly, all HTTPS responses flow through the HTTP proxy before being returned to the client application.

Provider Filtering Configurations

The settings that you can configure for the **WebProxy** provider are described in WebProxy Filters. This includes settings that you can specify for **Hostname Filter** and **Port Filter** in the **Advanced Settings-Microsoft-Pef-WebProxy** dialog. In addition, you have the option to specify **HTTPS Client Certificate** information in this dialog as well, which includes the capability to reuse client and server connections.

Event Keyword and Error Level Filtering

Similar to other system ETW providers that are registered on your system, you have the option to specify ETW event **Keyword** configurations and error **Level** filter settings for the **Microsoft-Pef-WebProxy** provider, to further refine the scope of events to be captured. The settings for these filters are available on the **ETW Core** tab of the **Advanced Settings – Microsoft-Pef-WebProxy** dialog. By default, all events are delivered by the **Microsoft-Pef-WebProxy** provider when no **Keywords** are selected. If you specify a particular **Keyword**, then the provider will deliver the events that are enabled by that **Keyword** only, if they occur in a trace. It also follows that multiple events are delivered when multiple **Keywords** are selected.

More Information

To learn more about configuring system ETW Providers, including event **Keyword** and error **Level** filters, see System ETW Provider Event Keyword/Level Settings.

To learn more about configuring PEF providers, including filters for the **Microsoft-PEF-WebProxy** provider, see the following topics:

Common Provider Configuration Settings Summary

WebProxy Filters

See Also

Built-In Trace Scenarios

Microsoft-PEF-NDIS-PacketCapture Provider

Microsoft-PEF-WFP-MessageProvider

WebProxy Filters

2 minutes to read

The **Microsoft-Pef-WebProxy** provider has two filters for which you can set values. The purpose of these filters is to isolate specific traffic and optimize Message Analyzer capture performance. An option is also provided that enables you to specify a client certificate file that Fiddler will provide in a given session to a site server that requires certification validation. After you open the **New Session** dialog from the Message Analyzer **File** menu and select the **Pre-Encryption for HTTPS Trace Scenario** from the **Select a trace scenario** drop-down in the dialog, you can configure the following settings on the **Provider** tab of the **Advanced Settings – Microsoft-Pef-WebProxy** dialog (access by clicking the **Microsoft-Pef-WebProxy** provider **Configure** link to the right of the provider **Id** in the **New Session** dialog):

- **Hostname Filter** — enables you to filter HTTP messages to and from a specific host, by specifying the host name in a format similar to the following:

  www.xxxxx.com

  This filter restricts the capture of HTTP operations initiated by the client browser to requests and responses to and from a single web host. You might set such a filter when troubleshooting communications from a poorly performing web server in cases where you need to isolate messages sent to and received from that server only.

- **Port Filter** — enables you to filter traffic to a specific port, such as 80, 443, 8080, and so on, by specifying a port value in integer format, similar to the following:

  443

  You might use such a filter to limit the traffic you capture to HTTPS operations only, so you can isolate and troubleshoot both requests from the client browser and responses from the server that were intended to be encrypted.

  > **NOTE**
  > For maximum positive impact on HTTP capture performance, you should set both of the specified filters to appropriate values.

- **HTTPS Client Certificate** — enables you to specify a certificate file (*.cer) that Fiddler will provide in a given session to a server that requires certification validation. This is because in some scenarios, Fiddler must act as the proxy that provides the security certificate. You will need to specify the path to the certificate file in the **HTTPS Client Certificate** text box.

Other settings that you can specify on the **Provider** tab of the **Advanced Settings – Microsoft-Pef-WebProxy** dialog consist of the following:

- **Reuse Client Connections** — given that Fiddler is a proxy, when it creates HTTP client connections on behalf of Internet Explorer or another browser, it will keep these connections alive for reuse to more speedily expedite client HTTP requests, rather than keep creating new ones and incurring the accompanying degradation to performance.

- **Reuse Server Connections** — the same as the above, only as applied to HTTP server connections.

WebProxy Provider Manifest

2 minutes to read

The manifest for the **Microsoft-Pef-WebProxy** provider defines its contract with an ETW Controller in terms of the events that are issued by the provider, which are in turn consumed by Message Analyzer. The **Microsoft-Pef-WebProxy** provider and its manifest are installed and registered on your machine when you install Message Analyzer. You can find the **Microsoft-Pef-WebProxy** provider manifest file in the following location:

```
C:\Windows\System32\Microsoft-Pef-WebProxy.man
```

In this file you can view the event definitions, tasks, opcodes, keyword bitmask and level specifiers, a channel definition, and template information.

More Information

To learn more about the contents of provider manifests, see Event Manifest.

To learn more about how Message Analyzer uses provider manifests, see Understanding Event Parsing with a Provider Manifest.

To learn more about obtaining a provider manifest if you are missing one, see Generating a Provider Manifest.

Microsoft-Windows-NDIS-PacketCapture Provider

5 minutes to read

The **Microsoft-Windows-NDIS-PacketCapture** provider works with several **Trace Scenarios** that are optimized to capture traffic on either a local or remote computer. The **Microsoft-Windows-NDIS-PacketCapture** provider is instrumented to support these capture scenarios on computers running the Windows 8.1, Windows Server 2012 R2, Windows 10, or later operating system. This means that you can use any **Trace Scenario** that includes this provider to target the capture of traffic from any local or remote host, as long as they are running one of these operating systems. You can even target local and remote traffic at the same time, as long as you specify the local host and target remote computer names in the **Target Computers** list in the **New Session** dialog during Live Trace Session configuration. However, you cannot successfully target any computer for remote capture if it is running a down-level operating system such as Windows 7, Windows 8, or Windows Server 2012. The **Microsoft-Windows-NDIS-PacketCapture** is also instrumented to work with the ETW infrastructure, which provides the mechanisms for controlling ETW Sessions, buffering captured data, and delivering events. Because the **Microsoft-Windows-NDIS-PacketCapture** filter driver works with this infrastructure, it can deliver the frames it captures at the Data Link Layer as ETW events.

> **NOTE**
> Message Analyzer uses Windows Management Instrumentation (WMI) remoting facilities for capturing data on remote computers.

Provider Functions

Message Analyzer provides several built-in **Trace Scenarios** that use the **Microsoft-Windows-NDIS-PacketCapture** provider, in which there are special features that enable you to create unique capture configurations. For example, in the **Microsoft-Windows-NDIS-PacketCapture** provider configuration, you can specify **ETW Core** provider settings, interface selection, and advanced filtering configurations when capturing traffic on local or remote hosts, virtual machine (VM) adapters, and Hyper-V-Switches, as described in Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

You can use the previously indicated configuration features when capturing messages with this provider in scenarios that are optimized for capture on local computers. However, note that because the **Microsoft-Windows-NDIS-PacketCapture** provider supports remote capture on computers running the Windows 8.1, Windows Server 2012 R2, Windows 10, and later operating systems, you can target hosts that are running any of these operating systems for remote capture with *any* **Trace Scenario** that uses the **Microsoft-Windows-NDIS-PacketCapture** provider, for example, even the **Local Network Interfaces** scenario.

Associated Trace Scenarios

The scenarios that use the **Microsoft-Windows-NDIS-PacketCapture** provider are described as follows. With these scenarios, you can capture message data at and above the Data Link Layer on computers that are running the Windows 8.1, Windows Server 2012 R2, Windows 10, or later operating system:

- **Local Network Interfaces** scenario — as previously indicated, you can apply the advanced provider settings and filtering configurations described in Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog to local traffic captures; however, some settings will have a different meaning, such as the **All Layers** setting, depending on whether you are capturing from the local host versus a VM that is serviced by a local Hyper-V-Switch. For more information about advanced filtering, see Configuring

Host Adapter and Hyper-V-Switch Filters.

- **Remote Network Interfaces** scenario — for more information about the requirements for capturing remote traffic, see Configuring a Remote Capture. Use this scenario to capture traffic from a remote host or virtual machine that is serviced by a Hyper-V-Switch.
- **Remote Network Interfaces with Drop Information** scenario — also includes several other ETW Providers that are configured with **Keyword** filters to provide event information that identifies dropped packets. Use this scenario to capture traffic from a remote host or virtual machine (VM) that is serviced by a Hyper-V-Switch, when you are interested in determining whether packets are being dropped, possibly by the NDIS stack of a host adapter or the extension layers of a Hyper-V-Switch that is servicing one or more VMs.
- **VPN** scenario — also includes several other ETW Providers which specify **Keyword** configurations that optimize the capture configuration for exposing events associated with VPN communications. Use this scenario to troubleshoot VPN-related issues.
- **Wired Local Area Network** scenario — also includes several other ETW Providers that specify **Keyword** configurations that optimize the capture configuration to expose Windows component events associated with the LAN. Use this scenario to troubleshoot a wired LAN and to expose operating system issues.
- **Wireless Local Area Network** scenario — also includes several other ETW Providers that specify **Keyword** configurations that optimize the capture configuration to expose Windows component events associated with the LAN. Use this scenario to troubleshoot a wireless LAN and to expose operating system issues.

Provider Filtering Configurations

The special filter types that you can configure in the **Advanced Settings** dialog for the **Microsoft-Windows-NDIS-PacketCapture** provider consist of the following:

- Adapter filters
- Layer filters
- Truncation filters
- Direction filters
- EtherType filters
- IP Protocol Number filters
- MAC Address filters
- IP Address filters

More Information
To learn more about configuring these special filters, see the Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog topic, which also includes a figure that shows the **Advanced Settings** dialog for this provider.

Event Keyword and Error Level Filtering

Similar to other system ETW providers that are registered on your system, you have the option to specify ETW event **Keyword** configurations and error **Level** filter settings for the **Microsoft-Windows-NDIS-PacketCapture** provider, to further refine the scope of events to be captured. The settings for these filters are

available on the **ETW Core** tab of the **Advanced Settings – Microsoft-Windows-NDIS-PacketCapture** dialog. By default, all events are delivered by the **Microsoft-Windows-NDIS-PacketCapture** provider when no **Keywords** are selected. If you specify any particular **Keyword**, then the provider will deliver the events that are enabled by that **Keyword** only, if they occur in a trace. It also follows that the provider will deliver multiple events when you select multiple **Keywords**, if they occur in a trace.

More Information
To learn more about configuring ETW Providers, including event **Keyword** and error **Level** filters, see System ETW Provider Event Keyword/Level Settings.

Capturing Data in P-Mode

To enable capturing messages in Promiscuous Mode with the **Microsoft-Windows-NDIS-PacketCapture** provider, you can select any adapter that supports P-Mode captures in the Interface Selection section of the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog. You can perform such captures in P-Mode with the use of this provider on the local computer or on a specified remote computer.

More Information
To learn more about capturing data in P-Mode, see Capturing Remotely in Promiscuous Mode.

See Also

Built-In Trace Scenarios

Understanding Event Parsing with a Provider Manifest

2 minutes to read

Message Analyzer traces are based on ETW Providers, which operate in an infrastructure that enables you to capture ETW events on your system. For Message Analyzer to be able to parse ETW messages from a particular provider, it must have a manifest for that provider. The manifest is a configuration file that defines the schema or format of the data that is delivered by the provider. When starting a Live Trace Session that is configured to use a specific system ETW provider, Message Analyzer typically obtains the required system manifest from the .exe or .dll of the registered provider component and builds an OPN description (parser) to represent its messages. This OPN description is then compiled and temporarily loaded into the Runtime so it can parse the ETW messages. If a trace file is saved and then transferred to another system, that system may not have the same provider components, component versions, or an appropriate provider manifest that is needed to enable parsing of the ETW messages.

This issue can be resolved in either of the following ways:

- The manifest information must be saved with the trace.
- The manifest file must be manually generated and stored on the new system.

> **NOTE**
> On computers that are running the Windows 10 operating system, Message Analyzer can now use dynamic messages generated from raw ETW events to enable the Runtime to parse ETL files with no manifest. This feature accommodates the new ETL file format on Windows 10, where the message format definitions are self-contained. In this scenario, an OPN description is not required.

However, on computers that are running an operating system that is earlier than Windows 10, you might have an event trace log (ETL) that utilized an ETW provider that is not registered in the particular system on which you are running Message Analyzer. As a result, Message Analyzer will not have access to the provider manifest and will therefore be unable to fully parse messages from that log. In this case, Message Analyzer will provide a simple level of parsing that produces messages in a general format. However, if the manifest was previously included with the log that contains the data you are loading, Message Analyzer will be able to fully parse the messages in the ETL file.

You might also save trace data on a source computer that will be further processed on other destination systems where the provider versions are unknown. If this is the case, Message Analyzer accommodates this situation by automatically saving trace data with the manifests of the underlying provider/s that were used in the trace. This ensures that Message Analyzer will be able to parse the ETW message data on the destination computer. Likewise, if you are loading data from an ETL file into Message Analyzer and you suspect that an unknown provider configuration was used to capture/log the data, you might need to generate a manifest for the log to ensure that Message Analyzer can parse its messages.

Obtaining a Provider Manifest

If you need to obtain a provider manifest for parsing ETW messages, see Generating a Provider Manifest.

See Also

Integrating Event Tracing
Locating Supported Input Data File Types

Generating a Provider Manifest

2 minutes to read

You might need to generate a provider manifest to ensure that Message Analyzer can parse ETW messages on a destination computer if the ETW events were captured on another source computer and the destination computer has no corresponding manifest to understand the data format. You might also have a destination computer with a manifest that is out-of-date with respect to the source computer manifest that was included in an ETL file. To accommodate these requirements, you can generate and save a manifest in the following ways:

- **Automatically** — requires that Message Analyzer is installed on the source computer. For this method, you configure a Live Trace Session and specify one or more ETW providers and then perform a capture of ETW events. When you save a trace that used one or more ETW providers to capture live data, Message Analyzer automatically attaches the manifest for each ETW provider in use to the trace.

  > **NOTE**
  >
  > Manifests for the **Microsoft-PEF-NDIS-PacketCapture**, **Microsoft-PEF-WFP-MessageProvider**, and **PEF-WebProxy** providers are installed on your system by default when you install Message Analyzer. Also, during Message Analyzer installation, all the system ETW providers that are registered on your system are enumerated and their manifests are automatically obtained.

- **Manually** — Message Analyzer is not installed on the source computer, or the manifest on the destination computer is out of date. In these cases, you must employ standard external tools to perform the trace and generate a manifest for the provider/s in use. For example, to capture the ETW events, you might use a command-line tool such as **netsh** or **logman**. When you are ready to save the manifest for the ETW trace, or if you already have an ETL file without a manifest, you can run the command-line tool *Tracerpt* with the following command string to generate the manifest file with a .man extension:

  ```
  Tracerpt -l [ETWTrace file | *.etl file] -export [filename.man]
  ```

  This command exports the schema for the ETW events that were captured and saved to the specified event trace file or log (*.etl). For more information about Tracerpt commands, specify the following help switch:

  ```
  Tracerpt /?
  ```

  > **TIP**
  >
  > You might be able to obtain a manifest from InsightWeb as well, if you have access to internal Microsoft sites.

> **IMPORTANT**
>
> As an alternative to the previous methods, you may be able to use the following tool to create a manifest:
>
> - A new Visual Studio extension known as the Microsoft EventRegister Tool is available from a Nuget package.
>
>   This extension includes an EventSource class and other related classes in the Microsoft.Diagnostics.Tracing namespace, which enable you to write to the Event Log and generate a manifest at build time. For additional background information, see Getting Started with Semantic Logging

At this point, you should have one of the following from the source machine:

- A .matu or .matp file with the manifest information automatically attached.
- A trace file or *.etl log with a separate *.man file.

If you have a .matu or .matp file with the required manifest attached, you can simply load the file and Message Analyzer should be able to parse the messages. If you manually generated the manifest for an ETL file, you will need to place the manifest <filename.man> in the following folder prior to loading the log data into Message Analyzer:

```
%localappdata%\Microsoft\MessageAnalyzer\OPNAndConfiguration\EtwManifests\
```

> **NOTE**
> In a future release, Message Analyzer may have the capability to import provider manifests directly through an OPN adapter, depending on demand.

See Also

Understanding Event Parsing with a Provider Manifest

Configuring a Live Trace Session

6 minutes to read

To capture data live with Message Analyzer, you can either start a basic Live Trace Session instantly with a single click, as described in Performing a Live Capture, or you can create and configure a Live Trace Session from the **New Session** dialog. If you choose the latter, you will need to specify a capture configuration of choice that involves performing one or more of the following tasks from the **New Session** dialog:

- Selecting a built-in **Trace Scenario** or a custom scenario of your own design from the **Select Scenario** drop-down list, or selecting a system ETW Provider in lieu of a **Trace Scenario**. This is a required task.
- Specifying target computers for data capture in the **Edit Target Computers** dialog, or accepting the default Local host.
- Optionally adding other system ETW Providers to the provider list of a selected **Trace Scenario** by choosing them from the **Add Providers** drop-down list, to customize the scope of data capture.
- Enhancing message provider configurations with driver-level filtering, by clicking the **Configure** link in the **ETW Providers** list and creating a filter in the **Advanced Settings** dialog.
- Setting a **Parsing Level** from the **Parsing Level** drop-down list in the **New Session** dialog to limit the stack level to which Message Analyzer will parse, for performance improvements and focusing trace results on specific messages.
- Selecting a built-in **Session Filter** from the centralized **Library** or creating a custom Filter Expression of your own design in the **Session Filter** text box, to create focus on specific data you want to analyze while limiting the display of irrelevant data.
- Choosing a data viewer from the **Start With** drop-down list, or accepting the default viewer setting.

This section describes how to perform these tasks by using the integrated features of Message Analyzer. In the discussions that follow, see the following figure for the location of referenced features.

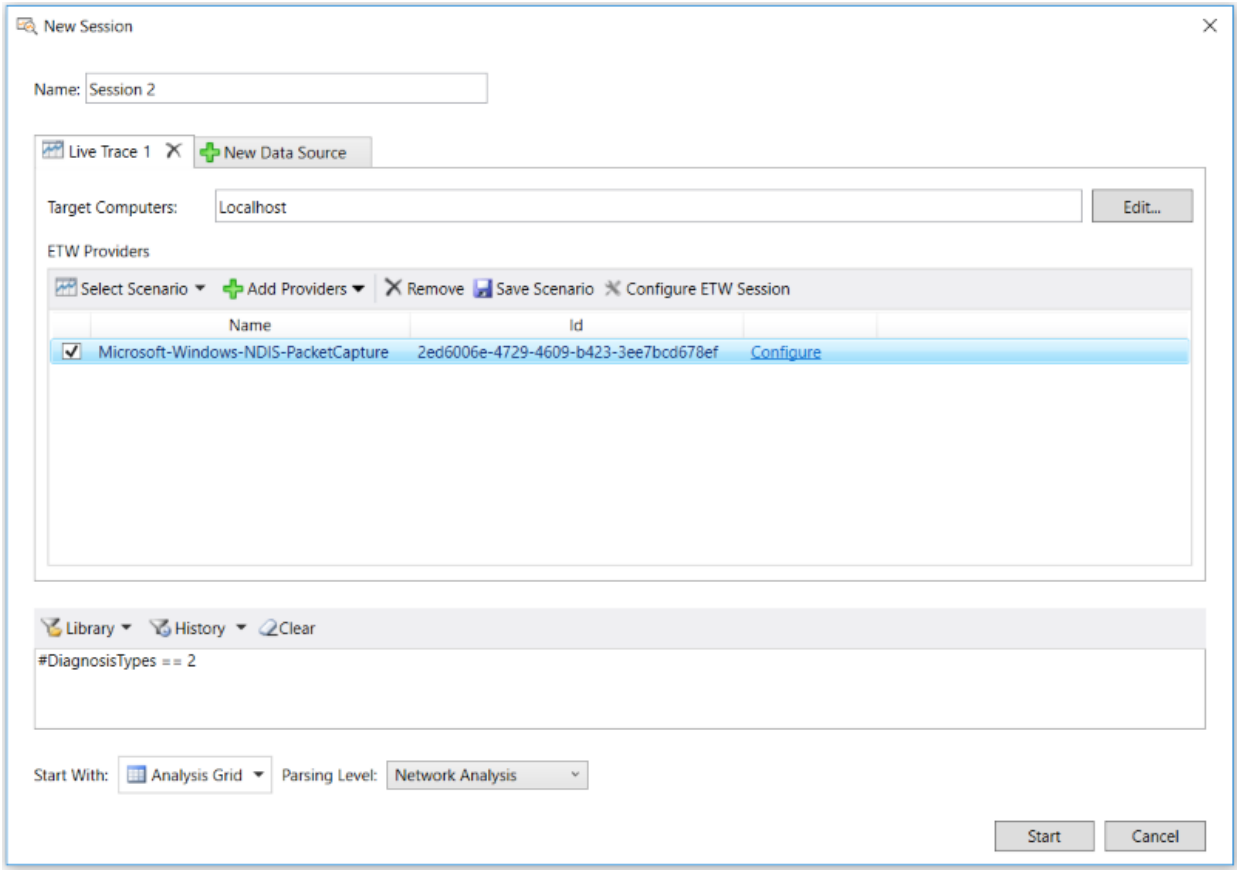

Figure 20: Live Trace Session configuration

Live Trace Session Workflow Overview

The following steps are an overview of the workflow that you can follow when configuring a Live Trace Session:

1. Begin new session configuration by clicking the **New Session** button on the Message Analyzer **Start Page** to display the **New Session** dialog, from where you specify a **Live Trace** as the data source, as described in Starting a Message Analyzer Session.

2. Click the **Live Trace** button in the **New Session** dialog to display the configuration features that you can use for your Live Trace Session.

3. Accept the default **Localhost** specification in the **Target Computers** text box of the **New Session** dialog, or specify connection information for one or more target computers from which to capture data in the **Edit Target Computers** dialog that is accessible by clicking the **Edit...** button on the **Live Trace** tab, as described in Configuring a Remote Capture.

4. Select a built-in **Trace Scenario** to focus on messages at a particular stack level or messages from a device, application, or target computer, as described in Selecting a Trace Scenario; you can also select a custom **Trace Scenario** template that you created, as described in Using a Custom Trace Scenario Template.

5. Optionally, select one or more system ETW Providers from the **Add System Providers** dialog, which is accessible from the **Add Providers** drop-down list on the **ETW Providers** toolbar of the **New Session** dialog, to enhance the scope of the data capture, as described in Adding a System ETW Provider.

6. Modify settings that can impact the performance and focus of message providers, which can include specifying **Fast Filters**, **WFP Layer Set** filters, adapter filters, NDIS stack filters, Hyper-V-Switch extension layer filters, packet filters, and event **Keyword** and **Level** filters, depending on the message provider/s in use, as described in Modifying Default Provider Settings.

7. Optionally, set the stack level to which Message Analyzer will parse, by specifying a **Parsing Level** in the

New Session dialog to improve performance and focus trace results, as described in Setting the Session Parsing Level, or use the default **Full** parsing specification.

8. Optionally, optimize the ETW buffer configuration and Flush Timer interval from the **ETW Session – Advanced Configuration** dialog, if you expect that packets will be dropped in your scenario, as related to the criteria described in Specifying Advanced ETW Session Configuration Settings.
9. Optionally, specify a **Session Filter** to focus on specific message data that meets the criteria of the Filter Expression that you select from the centralized **Library** on the toolbar above the **Session Filter** text box in the **New Session** dialog, or configure your own Filter Expression, as described in Working with Session Filters in a Live Trace Session.
10. Optionally, choose a data viewer from the **Start With** drop-down list in the **New Session** dialog for the display of Live Trace Session results, for example, the **Analysis Grid** viewer or the **Gannt** viewer, as described in Selecting a Session Data Viewer.
11. Optionally rename your scenario in the **Name** text box and click the **Save Scenario** button in the **New Session** dialog to save your scenario with the customized settings that you specified, so that you can run it again on demand, as described in Saving Trace Scenarios.
12. Click **Start** in the **New Session** dialog to begin capturing data with your configured Live Trace Session.

Trace Scenario Configuration Overview

When configuring a Live Trace Session, you can select from a number of built-in **Trace Scenarios** that contain various provider configurations that perform common or focused tasks, as indicated in the table in the Built-In Trace Scenarios section. You can also configure your Live Trace Session to use additional system ETW Providers that are accessible from the **Add System Provider** dialog, as described earlier. However, if you include a system ETW Provider in your Live Trace Session configuration, you should be familiar with how to use it, which includes how to configure it to write specific events that you are interested in capturing through **Keyword** selection. For example, for system ETW Providers that define error **Level** and event **Keyword** values, you can configure the error verbosity level and specific events to capture by selecting them on the **ETW Core** tab of the **Advanced Settings** dialog for any provider, to further refine the focus of the data retrieval action of the overall provider set in use.

You might be able to learn more about the events that a system ETW Provider writes if you can locate and review its manifest, which can define such information. This will ensure that you have a viable data capture configuration and understandable trace results. For more information about the **Keywords** that define the events that a provider captures, see Using System ETW Providers as an Input Source and Finding System ETW Provider Metadata.

Also, to narrow the focus of your Live Trace Session to retrieving specific message data of interest only, you can apply a **Session Filter** as previously described. However, you should note that applying a **Session Filter** can have an impact on parsing time.In addition, you can modify low-level provider settings and/or set a **Parsing Level** to improve performance and focus results.

Message Analyzer also provides you with the flexibility needed to ensure that your Live Trace Sessions are easily accessible and configurable, productive, and simple to save. Essentially, you can configure a Live Trace Session based on any of the following:

- A built-in **Trace Scenario**.
- A **Trace Scenario** that you previously created as a custom template.
- One or more Windows system ETW Providers.
- A combination of Microsoft-PEF and Windows system ETW Providers.

More Information

To learn more about creating your own **Trace Scenarios**, see Creating and Managing Custom Trace Scenarios.

Live Trace Session Configuration Features

The following subtopics provide further details about Message Analyzer's Live Trace Session configuration features along with various operations that Message Analyzer supports for live captures:

Selecting a Trace Scenario
Using a Custom Trace Scenario Template
Adding a System ETW Provider
Modifying Default Provider Settings
Selecting Data to Capture
Working with Session Filters in a Live Trace Session Specifying Advanced ETW Session Configuration Settings
Configuring a Remote Capture
Selecting a Session Data Viewer
Decrypting TLS and SSL Encrypted Data

Capturing the Data

When you are ready to capture data with Message Analyzer, see Performing a Live Capture to review several methods for starting a live capture.

See Also

Message Analyzer User Roles

Selecting a Trace Scenario

3 minutes to read

Message Analyzer provides the **Message Analyzer Trace Scenarios** asset collection Library which contains various built-in **Trace Scenarios**. These built-in scenarios typically expose common Message Analyzer usage scenarios that are optimized for capturing messages of specific types or at a particular network stack layer. They are all accessible by clicking the **Select Scenario** drop-down list on the **Live Trace** tab in the **New Session** dialog. You can also access any **Trace Scenario** that you have set to Favorite status in either of the following ways:

- Click the **Favorite Scenarios** item in the Message Analyzer **File** menu and then select a favorite scenario in the submenu that appears.
- Click a scenario in the **Favorite Scenarios** list on the Message Analyzer **Start Page**.

Note that whenever you select a **Trace Scenario** that has **Favorite** status, a Live Trace Session is immediately started with no additional session configuration required. This gives you a quick and convenient method for very quickly starting a Live Trace Session.

Setting Trace Scenarios as Favorites

You can set any **Trace Scenario** as a **Favorite Scenario** by clicking the **Edit Favorites** label on the Message Analyzer **Start Page**. When the **Trace Scenario** library displays, you can set a scenario as a Favorite by clicking the white star to the left of the scenario name, at which point, the star color changes to amber and the scenario is then added to the **Favorite Scenarios** list. You can undo the favorite status of any **Trace Scenario** by clicking the amber colored star, at which time the star color changes to white and the scenario is removed from the **Favorite Scenarios** list. You can also perform similar actions from the **Select Scenario** drop-down list in the **New Session** dialog.

Selecting a Built-In Trace Scenario

The built-in **Trace Scenarios** provide you with various predefined provider configurations in the following categories, as described in Built-In Trace Scenarios:

- **Network**
- **Device**
- **System**
- **File Sharing**

When you select a built-in **Trace Scenario** in any of these categories, the following items display in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog:

- The names of the providers used in the scenario and an associated GUID that identifies each one that appears in the **ETW Providers** list.
- A **Configure** link that enables you to access the **Advanced Settings** dialog that provides configuration settings on the following tabs:
  - **ETW Core** — contains event **Keyword** and **Level** filter settings for any ETW Provider associated with the scenario.
  - **Provider** — contains filter or other settings for any provider that is associated with the scenario, for

example, the **Microsoft-Windows-NDIS-PacketCapture** provider or any **Microsoft-PEF** provider.

> **NOTE**
>
> While the **Advanced Settings** dialog for **Microsoft-PEF** providers and the **Microsoft-Windows-NDIS-PacketCapture** provider will display both a **Provider** and **ETW Core** tab, most system ETW Providers that you add to Live Trace Session configuration from the **Add System Providers** dialog display an **ETW Core** tab only in the **Advanced Settings** dialog.

Selecting a Custom Configured Trace Scenario

If you select a custom created **Trace Scenario** template from the **My Items** category of the **Trace Scenarios** Library, all of the settings that you previously specified for the template are contained in your Live Trace Session configuration. This can include any combination of **Microsoft-PEF** providers, the **Microsoft-Windows-NDIS-PacketCapture** provider, system ETW Providers, event **Keyword** and error **Level** filter settings, **Fast Filters**, a **Session Filter**, a **Parsing Level**, advanced NDIS stack and Hyper-V-Switch extension layer filters, and so on. Note that you have the option to reconfigure and save any of these settings as required, before you run your custom **Trace Scenario** template again, as described in Using a Custom Trace Scenario Template.

After selecting a custom **Trace Scenario**, you have the option to either start capturing data immediately or to modify various provider settings and then start your Live Trace Session.

See Also

Built-In Trace Scenarios

Built-In Trace Scenarios

24 minutes to read

All Message Analyzer installations include a built-in set of predefined **Trace Scenarios** that together provide you with a large range of tracing functionality, applicability, and usefulness. These scenarios can help you get started very quickly with capturing and processing live data. For example, you can simply select a built-in **Trace Scenario** and then start your Live Trace Session with no additional configuration required. However, the quickest way to start a Live Trace Session is to click the **Start Local Trace** button on the **Start Page** to immediately begin capturing network messages on your local computer at the Data Link Layer and above by using the **Local Network Interfaces** scenario. Other quick start methods that you can use to immediately launch a Live Trace Session and start capturing live data consist of the following:

- Click one of the built-in **Trace Scenarios** in the **Favorite Scenarios** submenu, which is accessible from the Message Analyzer **File** menu.
- Click one of the built-in **Trace Scenarios** in the **Favorite Scenarios** list, which is accessible from the Message Analyzer **Start Page**.

Trace Scenarios Library

Message Analyzer maintains all items of the **Message Analyzer Trace Scenarios** asset collection in a Library that is accessible to the Message Analyzer Sharing Infrastructure, where you can auto-synchronize with collection updates that are pushed out by a Microsoft web service or manually download them as required with the use of the **Asset Manager**, which is accessible from the global Message Analyzer **Tools** menu. The built-in **Trace Scenarios** utilize different combinations of providers to achieve specific results that are useful in common network, component, and device troubleshooting scenarios. You also have the option to specify your own provider combinations, by adding more providers to a built-in **Trace Scenario** or by specifying chosen system ETW Providers for a custom **Trace Scenario** that you create.

Some examples of how you might customize ETW-instrumented message provider combinations consist of using the following:

- A single PEF provider, such as the **Microsoft-PEF-NDIS-PacketCapture** provider or the **Microsoft-PEF-WFP-MessageProvider** in a standalone configuration.
- The **Microsoft-Windows-NDIS-PacketCapture** provider in a standalone configuration.
- A PEF provider and a combination of one or more Windows system ETW providers.
- The **Microsoft-Windows-NDIS-PacketCapture** provider and a combination of one or more Windows system ETW providers.
- One or more Windows system ETW providers.

Operating System Dependencies

The built-in **Trace Scenarios** and the providers they utilize for capturing data are described in this section. Note that the **Trace Scenarios** that are available in the **Select Scenario** drop-down list in the **New Session** dialog are specific to the supported operating system you are running. For example, the **Local Network Interfaces Trace Scenario** in the **Network** category on computers running the Windows 7, Windows 8, or Windows Server 2012 operating system, uses the **Microsoft-PEF-NDIS-PacketCapture** provider; while computers running the Windows 8.1, Windows Server 2012 R2, Windows 10, or later operating system use

the **Microsoft-Windows-NDIS-PacketCapture** provider in the **Local Network Interfaces Trace Scenario**. In Message Analyzer, when a **Trace Scenario** has an operating system dependency, it is specified as part of the scenario name in the **Select Scenario** drop-down list or it is included in the scenario description. In any case, the **Select Scenario** drop-down list in the **New Session** dialog will never contain any **Trace Scenarios** that do not apply to the supported operating system your computer is running.

Whenever you select a scenario in the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, the providers that are included in each scenario display in the **ETW Providers** list, along with their **Ids** or globally unique identifiers (GUIDs). A short description of each **Trace Scenario** in this user Library is also included below the scenario name, and when there are environment differences, the operating system that supports the scenario is typically specified.

Configuring Remote vs Local Traces

There are differences in the way you can configure the **Microsoft-Windows-NDIS-PacketCapture** and **Microsoft-PEF-NDIS-PacketCapture** providers prior to running a trace, as follows:

- **Remote trace scenarios with the Microsoft-Windows-NDIS-PacketCapture** provider — in remote scenarios that use this provider, you can specify the *remote* host adapters and/or virtual machine (VM) adapters from which to capture messages, the manner in which packets traverse the NDIS stack layers or Hyper-V-Switch extension layers on such remote adapters, respectively, and various unique filters such as **Truncation**, **EtherTypes**, and **IP Protocol Numbers**. You can configure these settings from the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** provider dialog, as described in Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

  > **NOTE**
  > Although the **Microsoft-Windows-NDIS-PacketCapture** provider has remote capabilities, its ability to capture message data on local hosts is utilized in several other Message Analyzer **Trace Scenarios**, for example, when capturing messages at the Data Link Layer in the **Local Network Interfaces (Win 8.1 and later)** scenario.

- **Local trace scenarios with the Microsoft-PEF-NDIS-PacketCapture** provider — in local scenarios that use this provider, you can specify local adapters from which to capture messages, the direction in which to capture them, and you can create up to two logically-chained **Fast Filter Groups** that you can assign to any selected adapter. You can configure these settings from the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, as described in Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

> **NOTE**
> In Message Analyzer v1.3 and later, the **Microsoft-PEF-WFP-MessageProvider** has the capability to capture messages from remote computers that are running the Windows 10 operating system. You can capture this data in any **Trace Scenario** that uses this provider by starting your Live Trace Session with this scenario from any computer that is running the Windows 8.1, Windows Server R2, Windows 10, or later operating system.

Built-In Trace Scenario Descriptions

The built-in **Message Analyzer Trace Scenarios** asset collection items that are included with every Message Analyzer installation are described in the table that follows, along with a functional description and possible usage for each scenario.

Table 4. Message Analyzer Built-In Trace Scenarios

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
Network Category			
Local Network Interfaces (Win 8 and earlier) Capture local Link Layer traffic from NDIS. OS: Windows 7, Windows 8, and Windows Server 2012.	Microsoft-PEF-NDIS-PacketCapture	Provides the capability to capture local traffic on the indicated operating systems at the Data Link Layer (wire level), which is the lowest available chokepoint in the network stack. Also enables you to configure **Fast Filters** that do the following: - Target specific packet data. - Reduce CPU processing and consumption of resources by passing less data. - Prevent higher disk I/O overhead. - Improve speed by avoiding filtering at the parsing engine level. Note that packets captured at the Data Link Layer can be encrypted by a protocol such as IPsec, which obfuscates cleartext transmissions. Also, data obtained from the **Microsoft-PEF-NDIS-PacketCapture** provider can be noisy, especially on a wireless connection, because it captures broadcast and other traffic below the Network layer.	You might use the **Local Network Interfaces** scenario if you want to: - Capture raw data on the wire, such as Ethernet frames. - Specify the configuration of adapters from which to capture data. - Specify light-weight **Fast Filters** that enable you to locate messages that contain specified offset length patterns (OLP) or messages intended for specified target addresses. You can logically chain up to 3 **Fast Filters** within two separate filter **Groups** which you can then apply to selected adapters. **More Information** **To learn more** about configuring these settings, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
Local Network Interfaces (Win 8.1 and later) Capture local Link Layer traffic from NDIS. OS: Windows 8.1, Windows Server 2012 R2, and Windows 10.	Microsoft-Windows-NDIS-PacketCapture	Provides the capability to capture local traffic at the Data Link Layer on computers running the Windows 8.1, Windows Server 2012 R2, Windows 10, and later operating systems. Also enables you to capture local VM traffic on Windows Server 2008 R2 and Windows Server 2012 computers. Configuration features include special **Filters** that do the following: - Truncate packets to reduce bandwidth consumption. - Establish how packets traverse the NDIS filter stack. - Isolate Ethernet frames that contain IP packets such as IPv4 and IPv6. - Filter for and return only IP packets that have certain payloads, for example, TCP, UDP, or ICMP. - Filter traffic based on one or more specified MAC or IP addresses.	You might use the **Local Network Interfaces** scenario on a local computer running the Windows 8.1, Windows Server 2012 R2, Windows 10, or later operating system if you want to: - Capture raw data on the wire, such as Ethernet frames. - View only the packet headers for a particular protocol, through truncation. - Monitor NDIS filter layers to determine whether packets are being dropped. - Specify the direction in which packets traverse the NDIS filter layers, to isolate inbound or outbound traffic. - Filter for packets that are intended for a particular address or that contain specific payload types. **Tip:** You can use the **Microsoft-Windows-NDIS-PacketCapture** provider to capture traffic from not only local computers but remote computers as well. See Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog for more information.

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
Loopback and Unencrypted IPSEC Captures above IPv4/IPv6 layer using the Windows Filtering Platform. Exposes loopback traffic in two directions and unencrypted IPSEC traffic. OS: Windows 7 and later.	Microsoft-PEF-WFP-MessageProvider	The WFP capture system does the following in this scenario: - Captures loopback traffic and unencrypted IPsec traffic. - Supports data capture at various points in the Windows kernel TCP/IP stack, for example, above the IP/Network layer. - Logs structured packet data as ETW events for application protocol analysis and traffic monitoring. - Provides raw binary data. - Enables you to configure **Fast Filters** that focus the retrieval action of the **Microsoft-PEF-WFP-MessageProvider**. - Enables you to log discarded packet events. **Note:** If you select the **Select Discarded Packet Events** check box on the **Provider** tab in the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog, any **Fast Filter** or **WFP Layer Set** filter that you have also specified will not apply to packet events that are discarded.	You might use the **Loopback and Unencrypted IPSEC** scenario with the **Microsoft-PEF-WFP-MessageProvider** if you want to: - Focus on troubleshooting local application communication issues via loopback traffic, for example, between a SQL Server and a web server. - Focus on troubleshooting IP security issues by capturing and analyzing unencrypted IPSec traffic. - Isolate traffic above the IP/Network layer and minimize broadcast and other lower-layer noise. - Isolate inbound or outbound TCP/IP traffic for IPv4 and IPv6. - Specify light-weight port and address **Fast Filters** that enable you to select specific messages to capture. - Troubleshoot discarded packet issues. - Target a computer running the Windows 10 operating system for remote capture in a Message Analyzer v1.3 or later installation, from a computer that is running the Windows 8.1, Windows Server R2, or Windows 10 operating system. **Tip:** You can also use the **New-PefTargetHost** PowerShell cmdlet to capture traffic on a remote Windows 10 computer with the **Microsoft-PEF-WFP-MessageProvider**. **For more information**, see Automating Tracing Functions with PowerShell.

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
Pre-Encryption for HTTPS Capture HTTPS client-side unencrypted traffic by using the **Web Proxy**-Fiddler provider.	Microsoft-PEF-WebProxy	Provides the ability to capture Application Layer/HTTP client-side browser traffic prior to encryption. The **Pre-Encryption for HTTPS Trace Scenario** does not capture data from lower layers, such as the Transport layer or below. As a result, you may not capture all HTTP traffic of interest unless you run a **Loopback and Unencrypted IPSEC** or **Local Network Interfaces** trace. Note that the **Microsoft-Pef-WebProxy** provider will not capture traffic to and from a web browser unless you configure **Internet options** to use a proxy server for the LAN. **Important:** To use the **Microsoft-Pef-WebProxy** provider, you must have the Fiddler library from Telerik installed. If you have not already installed this library, you can download it here. For more information, see Microsoft-PEF-WebProxy Provider.	You might use the **Pre-Encryption for HTTPS** scenario if you need to: - Capture all HTTP traffic to and from a web browser in unencrypted format. - Troubleshoot Web server and client performance issues. - Filter HTTP traffic based on a hostname URL or a particular port number, such as 80 or 443. - View various sets of HTTP statistics, such as the number of requests and responses, reason phrases, status codes, IDs, host URIs, ports, query strings, server response times, and so on. - Specify a security certificate to capture HTTPS traffic on a particular site where it is necessary for Fiddler to provide such a certificate. - Configure Fiddler to reuse client and/or server connections for performance improvements.
Local Loopback Network Capture loopback network traffic that references the loopback addresses of 127.0.0.1 and ::1. If the traffic uses one of the local IP addresses, the scenario should be updated to include that address. Display addresses with the `IPConfig /all` command.	Microsoft-PEF-WFP-MessageProvider	Passes only loopback traffic that uses the IPv4 and IPv6 loopback addresses. Will also include loopback traffic that uses a local IP address if you specify a **Fast Filter** that contains that address.	The provider configuration for this scenario, which includes the use of the **Advanced Settings – Microsoft-Pef-WFP-MessageProvider** dialog, enables you to do the following: - Focus on troubleshooting local application communication issues via loopback traffic, for example, between a SQL Server and a web server. - Focus on inbound loopback traffic only for IPv4 and IPv6.

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
Network Tunnel Traffic and Unencrypted IPSEC Capture network traffic in the VPN/DirectAccess tunnel by using the Microsoft-PEF-WFP-MessageProvider. Also capture unencrypted IPSEC traffic.	Microsoft-PEF-WFP-MessageProvider	In this scenario, the Microsoft-PEF-WFP-MessageProvider captures VPN, Direct Access, and IPSEC traffic. You can also use this scenario to remove loopback traffic. However, you must manually specify **Fast Filters** for IPv4 and IPv6 to remove the loopback traffic, for example, specify `!127.0.0.1` for the IPv4 filter and `!::1` for the IPv6 filter. You can also realize improved performance in this scenario because it excludes traffic from the Network Layer and below.	The provider configuration for this scenario, which includes the use of the **Advanced Settings – Microsoft-Pef-WFP-MessageProvider** dialog, enables you to do the following: - Focus on troubleshooting network tunnel traffic. - Focus on troubleshooting IP security issues by capturing unencrypted IPSEC traffic. - Isolate traffic above the IP/Network layer, and minimize broadcast and other lower-layer noise. - Isolate inbound or outbound TCP/IP traffic for IPv4 and IPv6. - Specify light-weight port and address **Fast Filters** that enable you to select specific messages to capture. - Target a computer running the Windows 10 operating system for remote capture in a Message Analyzer v1.3 or later installation, from a computer that is running the Windows 8.1, Windows Server R2, or Windows 10 operating system.
Pre-Encrypted HTTPS Direct Captures HTTP directly from the Microsoft-Windows-WinInet provider for enhanced HTTPS troubleshooting. OS: Windows 10 and later.	Microsoft-Windows-WinInet-Capture	To capture HTTPS traffic with this scenario, you must be running Message Analyzer with elevated (Administrative) privileges. Use as an alternative to the **Microsoft-PEF-WebProxy**/Fiddler provider in the **Pre-Encryption for HTTPS** scenario, to allow better interoperability with web servers that require security certificates. This is because the **Microsoft-Windows-WinInet** provider captures HTTPS	The provider configuration for this scenario, which includes use of the **Advanced Settings – Microsoft-Windows-WinInet** dialog, enables you to do the following: - Capture data live with Message Analyzer in the **Pre-Encrypted HTTPS Direct** scenario, or use Windows in-box tools such as NetSh or Logman to capture HTTPS messages unencrypted and generate an ETL file that you can load and parse with

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
		request and response messages as events, at a place where encrypted packets are already decoded. The WinInet provider enables you to filter for events based on **Keyword** and **Level** configurations only, as specified in the **Advanced Settings – Microsoft-Windows-WinInet** dialog for this provider.	Message Analyzer. **Note:** The **Microsoft-Windows-WinInet** provider is included with every installation of the Windows 10 operating system. - Better accommodate some web server configurations that are less friendly to the **Windows-PEF-WebProxy**/Fiddler provider in the **Pre-Encryption for HTTPS** scenario. - Capture HTTPS traffic directly with computers that restrict Fiddler installations. Specify provider **Level** and/or **Keyword** filter settings, such as: - Event packets that contain **Critical**, **Error**, **Warning**, **Information**, or **Verbose** error details, as described in System ETW Provider Event Keyword/Level Settings. - **WININET_KEYWORD_SEND** — enables you to filter for outbound event packets. - **WININET_KEYWORD_RECEIVE** — enables you to filter for inbound event packets. - **WININET_KEYWORD_PII_PRESENT** — enables you to filter for event packets that possibly contain private information. - **WININET_KEYWORD_PACKET** — not used in this scenario. - **Microsoft-Windows-WinInet-Capture/Analytic** — not used in this scenario.
Remote Network Interfaces Remote capture on Link	Microsoft-Windows-NDIS-PacketCapture	Enables you to take advantage of the remote tracing capabilities of the	You might use the **Remote Network Interfaces** scenario if you want to:

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
Layer. OS: target machines with Windows 8.1, Windows Server 2012 R2, and Windows 10.	**Microsoft-Windows-NDIS-PacketCapture**	provider to capture traffic on a remote computer running the Windows 8.1, Windows Server 2012 R2, Windows 10, or later operating system at the Data Link Layer. With this provider, you can do the following: - Target specific remote hosts on which to capture traffic. - Specify the remote host adapters and/or VM adapters on which to capture data. - Create special packet and address filtering configurations.	- Capture raw Ethernet frames remotely. - Isolate traffic on a particular remote Windows 8.1, Windows Server 2012 R2, or Windows 10 host that you specify. - Isolate traffic on a specified host adapter or VM adapter on a remote Windows 8.1, Windows Server 2012 R2, or Windows 10 computer. **Important:** If you want to capture traffic from a specific remote VM, you will need to select the VM in the tree grid section of the **Advanced Settings** dialog and then specify a filter based on the VM **MAC Address** to isolate the data. Otherwise, you may capture Hyper-V-Switch traffic that is destined for all VMs that are serviced by the switch, given that a Hyper-V-Switch driver cannot of itself distinguish between VMs. - Specify packet traversal paths and filters for NDIS stack and Hyper-V-Switch extension layers, for example, when troubleshooting remotely dropped packets. - Perform special filtering that isolates message headers, messages that contain a particular type of payload, or messages intended for a particular physical or network address. **More Information** **To learn more** about configuring these settings, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog. **Tip:** You can capture remote traffic with the **Microsoft-Windows-NDIS-PacketCapture** provider in promiscuous

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
Remote Network Interfaces with Drop Information Remote capture on Link Layer including event data to indicate dropped messages. Truncation is set to 128 bytes. OS: target machines with Windows 8.1, Windows Server 2012 R2, and Windows 10.	Microsoft-Windows-NDIS-PacketCapture Microsoft-Windows-WFP Microsoft-Windows-NdisImPlatformEventProvider Microsoft-Windows-TCPIP Microsoft-Windows-Hyper-V-VmSwitch Microsoft-Windows-Qos-Pacer Microsoft-Windows-MsLbfoEventProvider Microsoft-Windows-Winsock-AFD	Enables you to take advantage of the remote tracing capabilities of the **Microsoft-Windows-NDIS-PacketCapture** provider to capture traffic on a remote computer running the Windows 8.1, Windows Server 2012 R2, Windows 10, or later operating system, in addition to also capturing dropped packet event information.	You might use the **Remote Network Interfaces with Drop Information** scenario if you want to: - Utilize the remote capabilities of the **Microsoft-Windows-NDIS-PacketCapture** provider, as previously described. - Log dropped packet events, the firewall rules that may have caused them to be dropped, and other drop event information.

mode. For further details, see Configuring a Remote Capture.

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
SASL LDAP pre-encryption Capture LDAP events that are already decoded. OS: target machines with Windows 7 or later.	Microsoft-Windows-LDAP-Client	Capture pre-encrypted LDAP frames and other information by using the LDAP client provider.	You can use the **SASL LDAP pre-encryption** scenario if you want to troubleshoot LDAP traffic to and from the Active Directory service. The configuration of this provider includes numerous preset event **Keyword** bitmask filters that will return events such as the following, if they are triggered during LDAP operations: - Search - Write - SSL - Bind - Serverdown - Connect - Bytes_received - Bytes_sent To review the full event configuration, click the ellipsis control (**...**) to the right of the **Keywords (Any)** text box to open the **ETW Keyword Filter Property** dialog with the auto-configured **Keyword** selection displayed. **More Information** **To learn more** about configuring event Keyword settings, see System ETW Provider Event Keyword/Level Settings.

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
VPN Troubleshoot VPN related issues. OS: Windows 8.1, Windows Server 2012 R2, and Windows 10.	Microsoft-Windows-NDIS-PacketCapture Microsoft-Windows-Ras-NdisWanPacketCapture Microsoft-Windows-NDIS Microsoft-Windows-IPSEC-SRV Microsoft-Windows-WFP Microsoft-Windows-TCPIP **Note:** Before running this scenario, deselect the **Microsoft-Windows-NDIS** provider in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog, since the **Microsoft-Windows-NDIS-PacketCapture** provider duplicates its functions.	Contains the **Windows-NDIS-PacketCapture** provider and other Windows system ETW providers that capture all Virtual Private Network (VPN) traffic on Windows 8.1, Windows Server 2012 R2, and Windows 10 computers.	You might use the **VPN** scenario if you want to: - Troubleshoot VPN issues by capturing Ethernet frames. - Utilize the configuration capabilities and settings that are described earlier in the **Local Network Interfaces (Win 8.1 and later)** scenario. **More Information** **To learn more** about configuring these settings, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog for the portions of this topic that apply to NDIS configuration and local tracing.
Wired Local Area Network (Win 8 and earlier) Troubleshoot LAN issues on Windows 7, Windows 8, and Windows Server 2012. Capture interface and component traffic to expose deep OS behavior. Similar to the "netsh trace start scenario=LAN" command.	Microsoft-PEF-NDIS-PacketCapture Microsoft-Windows-L2NACP Microsoft-Windows-Wired-Autoconfig Microsoft-Windows-EapHost Microsoft-Windows-OneX Microsoft-Windows-NDIS **Note:** Before running this scenario, you can uncheck the **Microsoft-Windows-NDIS** provider in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog, given that the **Microsoft-PEF-NDIS-PacketCapture** provider duplicates its functions.	Includes the **Microsoft-PEF-NDIS-PacketCapture** provider and other system ETW providers that write events related to the local/physical network connection.	You might use the **Wired Local Area Network (Win 8 and earlier)** scenario if you want to: - Troubleshoot connection issues related to network adapter configuration and VPNs. - Utilize the configuration capabilities and settings that are described earlier in the **Local Network Interfaces (Win 8 and earlier)** scenario. **More Information** **To learn more** about configuring these settings, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
Wired Local Area Network (Win 8.1 and later) Troubleshoot LAN issues for Windows 8.1, Windows Server 2012 R2, and Windows 10. Capture interface and component traffic to expose deep OS behavior. Similar to the "netsh trace start scenario=LAN" command.	Microsoft-Windows-NDIS-PacketCapture Microsoft-Windows-L2NACP Microsoft-Windows-Wired-Autoconfig Microsoft-Windows-EapHost Microsoft-Windows-OneX Microsoft-Windows-NDIS. **Note:** You can uncheck this provider in the **ETW Providers** list, given that the **Microsoft-Windows-NDIS-PacketCapture** provider has sufficient functionality for this scenario.	Includes the **Microsoft-Windows-NDIS-PacketCapture** provider and other system ETW providers that write events related to the local/physical network connection on Windows 8.1, Windows Server 2012 R2, and Windows 10 computers.	You might use the **Wired Local Area Network (Win 8.1 and later)** scenario if you want to: - Troubleshoot connection issues related to network adapter configuration and VPNs on a Windows 8.1, Windows Server 2012 R2, or Windows 10 computer. - Utilize the configuration capabilities and settings that are described in the **Local Network Interfaces (Win 8.1 and later)** scenario. **More Information** **To learn more** about configuring these settings, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog for the portions of this topic that apply to NDIS configuration and local tracing.
Wireless Local Area Network (Win 8 and earlier) Troubleshoot LAN issues for Windows 7, Windows 8, and Windows Server 2012. Capture interface and component traffic to expose deep OS behavior. Similar to the "netsh trace start scenario=WLAN" command.	Microsoft-PEF-NDIS-PacketCapture Microsoft-Windows-L2NACP Microsoft-Windows-EapHost Microsoft-Windows-OneX Microsoft-Windows-NDIS Microsoft-Windows-WLAN-Autoconfig Microsoft-Windows-NWifi Microsoft-Windows-VWifi **Note:** Before running this scenario, deselect the **Microsoft-Windows-NDIS** provider in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog, since the **Microsoft-PEF-NDIS-PacketCapture** provider duplicates its functions.	Includes the **Microsoft-PEF-NDIS-PacketCapture** provider and other system ETW providers that write events related to the wireless local area network connection.	You might use the **Wireless Local Area Network (Win 8 and earlier)** scenario if you want to: - Troubleshoot connection issues related to wireless network adapter configuration. - Utilize the configuration capabilities and settings that are described earlier in the **Local Network Interfaces (Win 8 and earlier)** scenario. **More Information** **To learn more** about configuring these settings, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
Wireless Local Area Network (Win 8.1 and later) Troubleshoot LAN issues on Windows 8.1, Windows Server 2012 R2, and Windows 10. Capture interface and component traffic to expose deep OS behavior. Similar to the "netsh trace start scenario=WLAN" command.	Microsoft-Windows-NDIS-PacketCapture Microsoft-Windows-L2NACP Microsoft-Windows-EapHost Microsoft-Windows-OneX Microsoft-Windows-NDIS Microsoft-Windows-WLAN-Autoconfig Microsoft-Windows-NWifi Microsoft-Windows-VWifi **Note:** Before running this scenario, deselect the **Microsoft-Windows-NDIS** provider in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog, since the **Microsoft-Windows-NDIS-PacketCapture** provider duplicates its functions.	Includes the **Microsoft-Windows-NDIS-PacketCapture** provider and other system ETW providers that write events related to the wireless local area network connection on Windows 8.1 or Windows Server 2012 R2 computers.	You might use the **Wireless Local Area Network (Win 8.1 and later)** scenario if you want to: - Troubleshoot connection issues related to wireless network adapter configuration. - Utilize the configuration capabilities and settings that are described earlier in the **Local Network Interfaces (Win 8.1 and later)** scenario. **More Information** **To learn more** about configuring these settings, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog for the portions of this topic that apply to NDIS configuration and local tracing.
Device Category			
Bluetooth (Win 8 and later) Troubleshoot Bluetooth issues.	Microsoft-Windows-BTH-BTHUSB	Contains Windows ETW providers that capture events related to Bluetooth devices.	You might use the **Bluetooth** scenario to troubleshoot a Bluetooth connection, pairing, and other issues, such as data display.
USB2 Troubleshoot USB 2 issues. OS: Any supported.	Microsoft-Windows-USB-USBPORT Microsoft-Windows-USB-USBHUB	Consists of two Windows providers that capture events related to USB2 devices.	You might use the **USB2** scenario to troubleshoot any device that is plugged into a USB2 port.
USB3 USB tracing for USB 3 host controllers (USB 2 or USB 3 devices). OS: Windows 8/Windows Server 2012 and later.	Microsoft-Windows-USB-USBXHCI Microsoft-Windows-USB-UCX Microsoft-Windows-USB-USBHUB3	Contains three Windows providers that capture events related to USB3 devices.	You might use the **USB3** scenario to troubleshoot any device that is plugged into a USB3 port.
System Category			

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
RPC Troubleshoot issues related to RPC framework.	Microsoft-Windows-RPC	Contains a single Windows provider that captures events from the remote procedure call (RPC) framework, including errors and other information (see the **Keyword** configuration for this provider).	You might use the **RPC** scenario to troubleshoot distributed programs that use RPC.
File Sharing Category			
SMB2 Client And Firewall Capture SMB2 client provider traffic with headers only, combined with the Microsoft-PEF-WFP-MessageProvider. Associate network traffic with SMB2 client traffic. OS: Windows 8, Windows Server 2012, or later.	Microsoft-Windows-SMBClient Microsoft-PEF-WFP-MessageProvider	Provides full SMB information in addition to message data above the IP/Network Layer with the Microsoft-PEF-WFP-MessageProvider.	You might use the **SMB2 Client And Firewall** scenario to support SMB2 client and firewall-level tracing.
SMB2 Client Full Payloads Capture SMB2 client provider traffic with the payload; exposes data being transferred in Reads and Writes. Also capture encrypted and DMA-transferred SMB traffic. OS: Windows 8, Windows Server 2012, or later.	Microsoft-Windows-SMBClient	Contains a single Windows provider that is extended for SMB client events.	You might use the **SMB2 Client Full Payloads** scenario to support tracing with SMB filtering so that you can see encrypted data from the SMB client. Provides better performance by filtering out data at the lower levels, such that only SMB packets are passed by the provider. **Tip:** The **ETW Core** configuration tab of the **Advanced Settings** dialog for all SMB providers in the **Windows 8 File Sharing** category exposes Keyword settings for additional filtering capabilities.
SMB2 Client Header Only Capture SMB2 client provider traffic without the payload; increases performance by capturing less data. Also capture encrypted and DMA-transferred SMB traffic. OS: Windows 8, Windows Server 2012, or later.	Microsoft-Windows-SMBClient	Contains a single Windows provider that is extended for SMB client events.	You might use the **SMB2 Client Header Only** scenario to support tracing with SMB filtering so that you can retrieve only the headers from packets sent by the SMB client. By capturing only the SMB headers, that is, without the data payload, this provider delivers significant performance improvements.

TRACE SCENARIO	PROVIDER NAMES	FUNCTIONAL DESCRIPTION	POSSIBLE USAGE CONFIGURATIONS
SMB2 Server Full Payloads Capture SMB2 server provider traffic with the payload; exposes data being transferred in Reads and Writes. Also capture encrypted and DMA-transferred SMB traffic. OS: Windows 8, Windows Server 2012, or later.	Microsoft-Windows-SMBServer	Contains a single Windows provider that is extended for SMB server events.	You might use the **SMB2 Server Full Payloads** scenario to support tracing with SMB filtering so that you can see encrypted data from the SMB server. Provides better performance by filtering out data at the lower levels, such that only SMB packets are passed by the provider.
SMB2 Server Header Only Capture SMB2 server provider traffic without the payload; increases performance by capturing less data. Also capture encrypted and DMA-transferred SMB traffic. OS: Windows 8, Windows Server 2012, or later.	Microsoft-Windows-SMBServer	Contains a single Windows provider that is extended for SMB server events.	You might use the **SMB2 Server Header Only** scenario to support tracing with SMB filtering so that you can retrieve only the headers from packets sent by the SMB server. By capturing only the SMB headers, that is, without the data payload, this provider delivers significant performance improvements.

More Information

To learn more about PEF provider capabilities, including capturing data with the network driver interface specification (NDIS) driver, see the PEF Message Providers topic.

To learn more about configuring provider settings, see Modifying Default Provider Settings.

To learn more about provider manifests, see Understanding Event Parsing with a Provider Manifest.

To learn more about managing the **Message Analyzer Trace Scenarios** asset collection, see Managing Trace Scenarios

See Also

Creating and Managing Custom Trace Scenarios

Using a Custom Trace Scenario Template

3 minutes to read

Instead of using a built-in **Trace Scenario** to start a Live Trace Session, you can start one by utilizing any custom **Trace Scenario** template that you created, as described in Designing a Trace Template. You create custom **Trace Scenario** templates so that you can store them for quick access to common capture functionality that is tailored to your environment and that you can use on a repetitive basis. You can display the provider configuration of any **Trace Scenario** template from the **Select Scenario** drop-down list on the **ETW Providers** toolbar of the **New Session** dialog by selecting the scenario during Live Trace Session configuration. Note that all custom **Trace Scenarios** that you create are stored in the **My Items** category in the specified drop-down list when you click the **Save Scenario** button on the **ETW Providers** toolbar.

To create a **Trace Scenario** template, you will need to start by first selecting an existing **Trace Scenario** and then modifying it as required, as described in Configuring a Live Trace Session. Thereafter, you can run your custom **Trace Scenario** template as you would any of the built-in **Trace Scenarios**, for example, the **Local Network Interfaces** scenario.

Running and Handling Trace Scenario Templates

Message Analyzer provides you with some flexibility when running and handling **Trace Scenario** templates. For example, after you select a trace template during Live Trace Session configuration, you can do any of the following:

- Run the template as-is to capture the specific type of message data that the template is configured to retrieve in your Live Trace Session. You can then save the results in one of the Message Analyzer native trace file formats, as described in Saving Session Data.
- Modify the template and then run it in a Live Trace Session to capture message data.
- Update the template and then save it without running it.
- Export the template in a specified asset collection, which can contain built-in **Trace Scenarios** and any scenario templates that you create, so that you can share them with others; either on a user-designated file share or through a user-configured feed in the Message Analyzer Sharing Infrastructure.

> **TIP**
> You can also import a **Trace Scenario** asset collection that others have shared to a designated location through the Message Analyzer Sharing Infrastructure.

- Set a custom **Trace Scenario** template to Favorite status in the **Trace Scenario** user Library, so that you can instantly start a Live Trace Session that uses such a scenario. All **Favorite Scenarios** are available from the Message Analyzer **Start Page**, the global Message Analyzer **File** menu, and the global Message Analyzer tool bar. Note that if you set a custom **Trace Scenario** template to Favorite status, it will be highly accessible from numerous locations in the Message Analyzer user interface (UI). For information on setting favorites, see Selecting a Trace Scenario.

The distinction between a built-in **Trace Scenario** and a **Trace Scenario** template is minor. All **Trace Scenarios** that are accessible from the **Select Scenario** drop-down list on the **ETW Providers** toolbar of the **New Session** dialog are essentially trace configuration templates, because a template does not yet contain data.

When you create a particular trace configuration and click the **Save Scenario** button, you create your own *custom*

trace configuration template that you can run at any time, just like any built-in **Trace Scenario**. When you save such a **Trace Scenario** template as indicated, the template itself will contain no data. But similar to all **Trace Scenarios**, when you run a **Trace Scenario** template in a Live Trace Session that captures data, you can then save the Live Trace Session *results* in the Message Analyzer native trace file .matp or .cap format. This file will contain your actual message data, whether the initial Live Trace Session provider configuration was a result of selecting a custom **Trace Scenario** template of your own design, or the result of selecting one of the Message Analyzer built-in **Trace Scenarios**.

> **NOTE**
>
> You can run a custom **Trace Scenario** and save session results as many times as you want. If you maintain a common capture configuration in a particular **Trace Scenario**, you can create a baseline for comparing similar trace results from day to day.

More Information

To learn more about **Trace Scenario** templates, see Creating and Managing Custom Trace Scenarios.

To learn more about exporting and importing **Trace Scenarios**, including **Trace Scenario** templates, see Managing Trace Scenarios.

Creating and Managing Custom Trace Scenarios

2 minutes to read

This section discusses the advantages of developing **Trace Scenarios**, how to create one by configuring a custom trace template, and how to save it as a **Trace Scenario** that encapsulates common, predefined capture functionality that you can access and run at any time.

These discussions are covered in the following subtopics:

Trace Scenario Overview
Creating Custom Trace Configurations
Designing a Trace Template
Saving Trace Scenarios
Starting a Custom Trace Scenario
Managing Trace Scenarios

Trace Scenario Overview

2 minutes to read

You might need to have quick access to common trace configurations that you can use on a regular basis to capture messages. Message Analyzer anticipates this need by enabling you to design and save *trace templates* that contain specific providers and filters that you configure for specific purposes. To do this, you can use the configuration features of a Live Trace Session to design and then save the template as a **Trace Scenario**. Your trace template can encapsulate specific tracing functionality of your own design, with one or more predefined providers and optional filter configurations that you can either run as-is, or modify to further customize its functionality before running it or resaving it. After you save the trace template, it becomes part of the **Trace Scenarios** Library and you can then access it as you would any other **Trace Scenario**.

To be clear, a distinction should be made between the concept of a **Trace Scenario** and an actual Live Trace Session. Generally speaking, all **Trace Scenarios** are trace templates. Therefore any custom trace template of your own design, and even the predefined **Trace Scenarios** that are provided with Message Analyzer by default, contain no data, as they are simply *templates* for a Live Trace Session. At the time when you select a **Trace Scenario** and providers appear in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog, you have a predefined provider configuration and you can optionally configure filters to define the scope of messages that you capture, but you do not yet have any trace results. A **Trace Scenario** becomes a Live Trace Session after you start the trace and capture data with it. Thereafter, you can save the Live Trace Session results in one of the Message Analyzer native file formats.

You can run any **Trace Scenario** that you wish on demand, including any custom trace template that you saved as a **Trace Scenario** in the **Trace Scenarios** Library. In addition, all **Trace Scenario** Library items are part of the Message Analyzer Sharing Infrastructure. As a result, you and other team members have mutual shared access to custom **Trace Scenarios** that you create. To export or import one or more **Trace Scenarios** to or from a designated file share or other location, you can use the **Manage Trace Scenario** dialog, which is accessible by clicking the **Manage Trace Scenarios** button from the **Select Scenario** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog.

More Information

To learn more about the Message Analyzer Sharing Infrastructure, see the Sharing Infrastructure topic.

To learn more about the common **Manage <AssetType>** dialog, see Managing User Libraries.

Creating Custom Trace Configurations

2 minutes to read

Message Analyzer enables you to create your own **Trace Scenario** templates by customizing the properties of an existing scenario and then saving your customized version as a **Trace Scenario** template with a specified name and category, without actually starting a trace and collecting data in a Live Trace Session. When you save **Trace Scenario** templates, they become part of the **Trace Scenarios** Library where they represent a master trace configuration that can include one or more providers, **Fast Filters**, a **Session Filter**, ETW **Keyword** bitmask or error **Level** filters, or advanced filters that you specify in the **Trace Scenario** configuration prior to saving the template. You can also choose the data viewer that your **Trace Scenario** will use by default whenever you decide to run the template configuration in a Live Trace Session.

Some of the advantages of creating your own preconfigured **Trace Scenarios** are that you can do the following:

- Create any number of **Trace Scenarios** that serve as trace templates for common usage contexts, where you focus on retrieving unique message data in each context.
- Have the convenience of quick access to a **Trace Scenario** template that specifies a trace configuration you typically run on a consistent basis.
- Share **Trace Scenario** templates that might be useful to your team members or the larger Message Analyzer community.
- Obtain useful **Trace Scenarios** from other team members or community sources.

Designing a Trace Template

2 minutes to read

Designing a trace template consists of adding providers, filters, and other settings to the trace configuration, just as you would do when configuring any Live Trace Session prior to running it, only you save the trace template as a **Trace Scenario** in the **Trace Scenarios** Library instead of running it. From the **Select Scenario** drop-down in the **New Session** dialog, you can add a default **Trace Scenario** with a predefined provider configuration to your trace template design by selecting one in the **Trace Scenarios** Library. You can then modify the provider configuration to achieve specific results that you require. For example, you might create a filter configuration that retrieves specific data in a Live Trace Session where you run your scenario template. Moreover, you can do the following and more when configuring your own trace template:

- Specify driver-level **Fast Filter Groups** in the default **Local Network Interfaces (Win 8 and earlier) Trace Scenario** provider settings, by using the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog that is accessible by clicking the **Configure** link for the **Microsoft-PEF-NDIS-PacketCapture** provider in the **ETW Providers** list of the **New Session** dialog.
- Add a custom written **Session Filter** to your trace template or add one that is predefined from the centralized Filter Expression **Library**.
- Modify certain default **Trace Scenario** provider capture settings, such as specifying **WFP Layer Set** filters, WFP **Fast Filters**, or dropped packet logging for **Trace Scenarios** that use the **Microsoft-PEF-WFP-MessageProvider**.
- Specify NDIS stack or Hyper-V-Switch extension **Layer** and packet **Direction** filtering, packet **Truncation**, **EtherTypes**, and **IP protocol numbers**, and filter on **MAC addresses** and **IP addresses** in the **Remote Network Interfaces (Win 8.1 and later) Trace Scenario** that uses the **Microsoft-Windows-NDIS-PacketCapture** provider. These settings are available from the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog.
- Customize the trace template configuration by adding system ETW providers that enhance the scope of data capture.
- Modify or add an event **Keyword** bitmask filter and/or error **Level** settings in the system **ETW Provider** configuration, if the provider defines them.
- Use the **ETW Session – Advanced Configuration** dialog to specify advanced buffer configuration settings that are passed to the ETW Session Controller that will manage your Live Trace Session.

More Information

To learn more about modifying provider configurations, see Modifying Default Provider Settings.
To learn more about saving a trace template as a **Trace Scenario**, see Saving Trace Scenarios.
To learn more about creating **Fast Filters** for the **Local Network Interfaces (Win 8 or earlier)** scenario, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.
To learn more about how to configure advanced filters for **Remote Network Interfaces** scenarios, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.
To learn more about creating filter expressions, see Writing Filter Expressions.
To learn more about how to configure advanced settings for ETW sessions, see Specifying Advanced ETW Session Configuration Settings.

Saving Trace Scenarios

2 minutes to read

When you are finished designing a trace template, you can save it as a custom **Trace Scenario** by clicking the **Save Scenario** button in the **New Session** dialog, at which time the **Edit Trace Scenario** dialog will display. In this dialog, you can specify a **Name** for the custom **Trace Scenario**, a **Description**, and add it to an existing **Category** or specify a new one. When you finish specifying your save configuration, click the **Save** button in the **Edit Trace Scenario** dialog to save your trace template as a new **Trace Scenario** item in the **Trace Scenarios** Library. The customized **Trace Scenario** that you are saving will be added to the **Trace Scenario** Library and will be available thereafter for selection from the **My Items** category of the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog during Live Trace Session configuration.

As part of this Library, your **Trace Scenario** template is accessible through the Message Analyzer Sharing Infrastructure. This infrastructure enables you to share your custom **Trace Scenario** templates with other team members and for others to share theirs with you.

For **Trace Scenario** Library items, the entry point into the Message Analyzer Sharing Infrastructure is accessed by clicking the **Manage Trace Scenarios** item in the **Select Scenario** drop-down menu on the **Live Trace** tab of the **New Session** dialog. When you do, the **Manage Trace Scenario** dialog displays all the **Trace Scenarios** in your Library in various categories, consisting of the default **Trace Scenarios** that are shipped with Message Analyzer, any **Trace Scenarios** that you have downloaded from Microsoft or imported through the Message Analyzer sharing infrastructure, along with any custom **Trace Scenarios** that you created. These custom **Trace Scenarios** will also be available for sharing with others through the Message Analyzer sharing infrastructure, through use of the **Export** feature in the **Manage Trace Scenario** dialog.

> **NOTE**
>
> To save a custom **Trace Scenario**, follow the procedure specified in Creating and Saving a Customized Trace Scenario.

Starting a Custom Trace Scenario

2 minutes to read

After you have created a trace template and saved it as a **Trace Scenario** in the **Trace Scenarios** Library, you can access the scenario and run it as needed. To do so, you will load your custom **Trace Scenario** into a Live Trace Session configuration by clicking it in the Library. You can then start the session as you would any Live Trace Session, by clicking the **Start** button in the **New Session** dialog. Note that it is unnecessary to specify a data viewer in which to display your trace results prior to starting your session, as your custom scenario will use whichever data viewer was specified in the **Start With** drop-down list at the time you saved the scenario. However, note that you still have the option to select a different data viewer prior to starting your Live Trace Session.

Managing Trace Scenarios

5 minutes to read

Message Analyzer provides facilities for managing your local **Trace Scenarios** Library. This means you can create and modify **Trace Scenarios**, set them as **Favorites** (by clicking the star to the left of a scenario in the **Trace Scenarios** Library), and share them with others. The Message Analyzer Sharing Infrastructure enables you to choose **Trace Scenario** items in your local Library and share them with others, and you can also add **Trace Scenario** items to your local Library from others. To do this, Message Analyzer enables you to **Export** your **Trace Scenarios** to a designated location to make them accessible to other users on your team or to the larger Message Analyzer community; or you can **Import Trace Scenarios** from other users who make their scenarios available to you in the same manner. This includes any custom **Trace Scenario** templates that you or others have created. Designated share locations can include a user-defined SMB share, a Web server resource, or a custom subscriber feed in the Message Analyzer Sharing Infrastructure that you configure from the **Settings** tab in the **Asset Manager** dialog, which is accessible from the Message Analyzer global **Tools** menu. When you export one or more **Trace Scenarios**, they are saved to a designated location as an asset collection file (*.asset) so that others can import one or more items in the collection.

Manage Trace Scenarios Dialog

You can manage **Trace Scenarios** from the **Manage Trace Scenario** dialog, which displays when you click the **Manage Trace Scenarios** item in the **Select Scenario** drop-down list on the **ETW Providers** toolbar of the **New Session** dialog. From the **Manage Trace Scenario** dialog, you can export and import **Trace Scenarios**, in addition to performing other scenario management tasks, as described in the subtopics that follow.

Exporting Trace Scenarios

The **Manage Trace Scenarios** dialog enables you to select the items you want to export in an asset collection for sharing or backup. After you select one or more **Trace Scenarios** that you want to share with others, click the **Export** button on the toolbar of the **Manage Trace Scenario** dialog to display the **Save Library** dialog, from where you can specify **Title**, **Description**, and **Organization** information. After you click **Save** in the **Save Library** dialog, the **Select Library Location** dialog displays and enables you to navigate to the location where you want to save your **Trace Scenario** asset file, which is typically a remote share or other feed that you configure from the **Settings** tab of the **Asset Manager** dialog.

> **NOTE**
>
> In a future Message Analyzer release, the Sharing Infrastructure may support the full publishing process for Message Analyzer asset collections that are shared on user-configured feeds. Currently, you can download **Trace Scenario** asset collections from user feeds, but you cannot yet configure them to automatically synchronize with periodic updates, unless you perform the manual configuration process described in Manual Item Update Synchronization. Also, you can always synchronize to and download asset collection updates from the **Message Analyzer** feed on the **Downloads** tab of the **Asset Manager** dialog to obtain the latest asset collections, such as **Trace Scenarios**, **Filters**, **Color Rules**, **View Layouts**, and so on.

Importing Trace Scenarios

The **Manage Trace Scenarios** dialog enables you to import asset collections that others have shared to some predefined location. To perform this task, click the **Import** button on the toolbar of the **Manage Trace Scenarios** dialog to display the **Select Library to Open** dialog. From this dialog, you can navigate to the local or network location where asset files are stored for sharing. After you select a *.asset file, click the **Open** button in the **Select Library to Open** dialog to display the **Select Items to Import** dialog. From this dialog, you can specify the items

from the asset collection that you want to import, along with the **Category** in which to place the imported scenarios in your local **Trace Scenarios** Library, or you can accept the default **Category** location. Click **OK** to exit the dialog and post the **Trace Scenarios** to the Library **Category** that you specified.

> **NOTE**
>
> **Trace Scenario** items that you import will appear under the **My Items** top-level category in the subcategory name that you specified, after a Message Analyzer restart.

Other Management Tasks

The **Manage Trace Scenario** dialog also enables you to open the **Edit Trace Scenario** dialog by right-clicking any **Trace Scenario** in the **My Items** category of the **Trace Scenarios** library and then selecting the **Edit** item from the context menu that displays. From the **Edit Trace Scenario** dialog, you can modify certain metadata that is associated with any scenario before exporting it as an asset file, as follows:

- **Name** — change the name of a **Trace Scenario**.
- **Description** — modify the **Trace Scenario** description.
- **Category** — select a different **Category** in which to place the scenario, or you can create a new **Category**.

> **NOTE**
>
> You cannot edit or delete any of the predefined **Trace Scenarios** that are provided by default in your Message Analyzer installation. You can edit scenario information or delete scenarios under the top-level **My Items** category only. The **My Items** category is reserved as user space in the **Trace Scenarios** library.

You can also create a copy of any **Trace Scenario** in any category, for example to move a copy to the **My Items Category**, by right-clicking any **Trace Scenario** in the **Manage Trace Scenario** dialog and selecting the **Create a Copy** item in the context menu that displays. When you do this, the **Edit Trace Scenario** dialog displays, from where you can modify any of the previously indicated parameters associated with the scenario and then **Save** it. Lastly, if you want to remove a **Trace Scenario** from the **My Items** category of the Library, you can right-click the scenario and select the **Delete** item in the same context menu. Note that the **Delete** command in the right-click context menu is disabled for all of the predefined **Trace Scenarios**, which prevents the removal of any of the predefined scenarios in your Message Analyzer installation.

More Information

To learn more about the Message Analyzer Sharing Infrastructure and the common **Manage <AssetType>** dialog format that is used to manage various Message Analyzer Library asset collections, see the Sharing Infrastructure topic.

Adding a System ETW Provider

9 minutes to read

Message Analyzer enables you to expand or narrow the scope of the data that a Live Trace Session will capture by specifying a system ETW Provider configuration. For example, you can add a system ETW Provider to the provider configuration that a built-in **Trace Scenario** provides by default, to capture events from a particular Windows component that will appear as additional data in your trace results — assuming that such events were written by the component's ETW Provider during the trace. Alternatively, you can select only system ETW Providers to narrow the focus to capturing events from specific components only, but you will need to be familiar with the component providers and the **Keywords** that you can enable for capturing specific events of interest.

The system ETW Providers that are registered on your system write events that are issued by various Windows components that have been instrumented with ETW technology to write such events. These providers are accessible from the searchable **Available System Providers** list in the **Add System Providers** dialog, which is accessible by clicking the **Add Providers** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog during Live Trace Session configuration. Note that many of these providers are based upon managed object format (MOF) schemas to define their events for ETW.

When you are considering which ETW Provider to add to a Live Trace Session, your familiarity with event tracing can be of significant value, given that you are likely to already understand the functions of various ETW Providers. However, if your experience is limited, choosing a provider may be a little more challenging. Many ETW Provider names are somewhat cryptic and may not adequately describe the provider functionality, whereas others are more clearly named. For example, the type of data captured by the **Microsoft-Windows-Dhcp-Client** provider is, generally speaking, not too difficult to determine. Moreover, if you review the **Keywords** for this provider, you will learn that it captures related system events in addition to others that reflect response time and client operational or administrative events. You might also try this approach when assessing other ETW Providers for possible inclusion in a Live Trace Session, because formal documentation for the large number of system ETW Providers is limited on Windows computers.

> **NOTE**
> You may be able to determine more about the events that are written by the ETW Providers on your system by examining the event **Keyword** configurations of such providers. This information typically exists in the manifests for the ETW Providers, which you can expose in several different ways, as described in Finding System ETW Provider Metadata.

Selecting a System ETW Provider in Live Trace Session Configuration

Once you know which system ETW Providers are of interest for the type of tracing you want to perform, the process of selecting a system ETW Provider to add to the **ETW Providers** list in the **New Session** dialog is straightforward. However, because the system ETW Provider list is very long, Message Analyzer provides a search box that enables you to locate providers quickly by name or GUID. To search for a system ETW Provider to add to the Live Trace Session configuration, perform the following steps.

1. Open the **New Session** dialog by clicking the **New Session** button on the global Message Analyzer toolbar.
2. Click the **Live Trace** button in the **New Session** dialog and then click the **Add Providers** drop-down list on the toolbar in the **ETW Providers** pane of the **New Session** dialog.

3. In the drop-down list, select the **Add System Providers** item to display the **Add System Providers** dialog that contains a list of system ETW Providers that Message Analyzer enumerated on your local system during installation, as shown in the figure that follows.

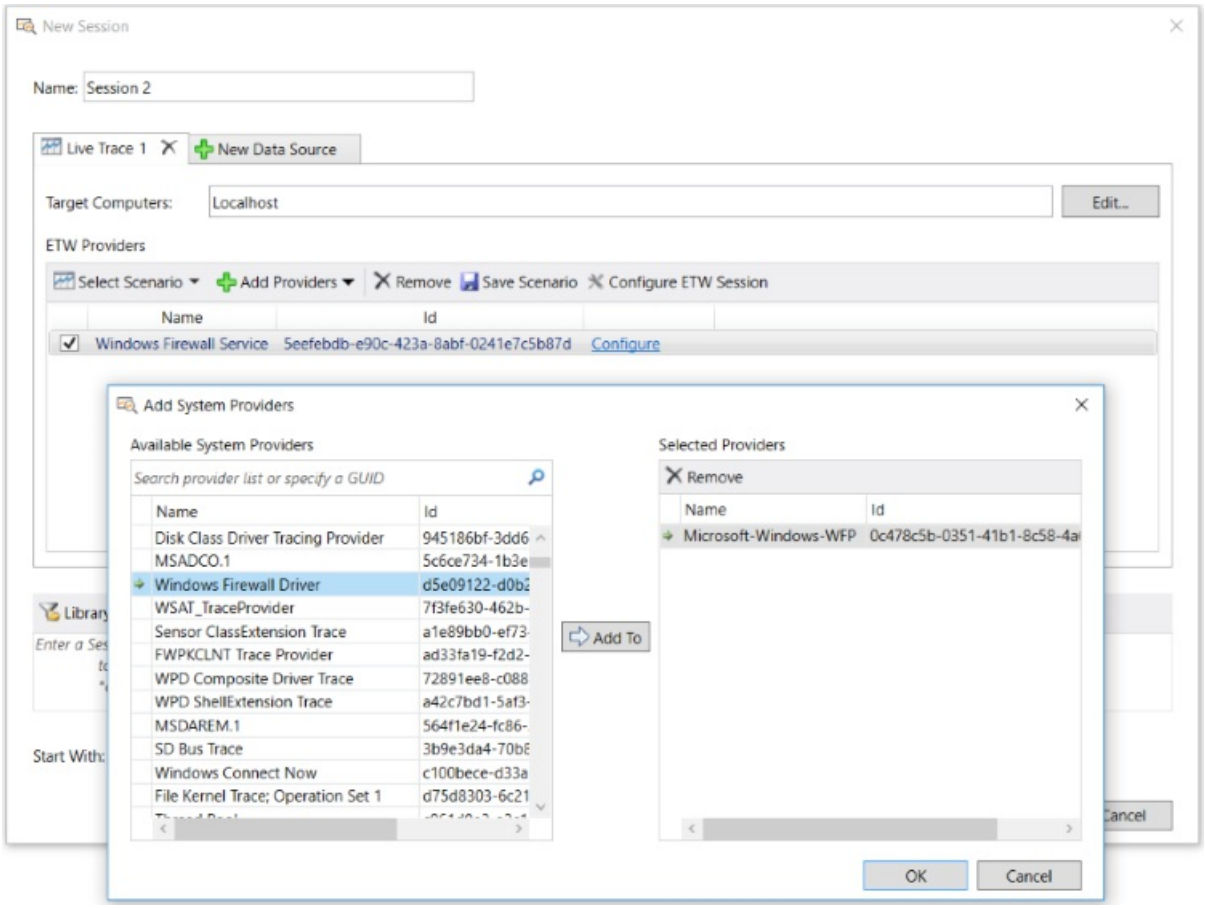

Figure 21: Specifying ETW Providers for Live Trace Session Configuration

4. Scroll down the **Available System Providers** list to locate a particular provider by name, or alternatively, type an ETW **System Provider** name or GUID in the search box to locate the provider of interest.

If you have a general idea of the type of provider you want to locate, you can enter text in the search box that partially reflects the provider name and Message Analyzer will return all provider names that contain the text you entered. For example, if you specify the text "dhcp", Message Analyzer will return a list of ETW Providers such as the **Microsoft-Windows-Dhcp-Client** and **Microsoft-Windows-DHCPv6-Client**.

5. After you locate the ETW Provider/s you want to add to the Live Trace Session configuration from the **Available System Providers** list of the **Add System Providers** dialog, highlight each one separately and then click the **Add To** button in the dialog to populate the **Selected Providers** list of the dialog.

6. After all the ETW Providers that you are adding to the Live Trace Session configuration are displayed in the **Selected Providers** list, click the **OK** button to exit the **Add System Providers** dialog.

At this time, the ETW providers that you are adding to the Live Trace Session configuration appear in the **ETW Providers** list of the **New Session** dialog.

7. Optionally modify the event **Keyword** and/or **Level** settings for any ETW Provider, as described in Configuring System ETW Providers.

IMPORTANT

You have the option to add a custom provider of your own to the **Available System Providers** list by specifying a **Guid** and a **Name** in the **Add Custom Provider** dialog. This dialog displays when you select the **Add Custom Provider** item in the **Add Providers** drop-down list on the toolbar of the **Live Trace** tab in the **New Session** dialog.

Configuring System ETW Providers

After you select a system ETW Provider as previously described and it displays in the **ETW Providers** list, you can access the configuration settings for the provider on the **ETW Core** tab of the **Advanced Settings** dialog. This dialog displays when you click the **Configure** link that appears immediately to the right of the **Id** for any provider that is listed in the **ETW Providers** list. If you want to further refine the focus of the provider's data retrieval action, you can modify the provider filtering configuration.

For example, you can specify event **Keyword** and **Level** filtering settings if the particular system ETW Provider defines such filters. To specify a **Level** filter from the **Advanced Settings** dialog, click the **Level** drop-down list and select an error **Level** that will cause the ETW Provider to return only the events that reflect the specified **Level** that you set. Note that **Level** filters are inclusive, as described in the "System ETW Provider Keyword and Level filter" Table in the System ETW Provider Event Keyword/Level Settings topic.

To enable a **Keyword** filter from the **Advanced Settings** dialog, click the ellipsis (**...**) to the right of the **Keywords (Any)** or **Keywords (All)** text boxes to display the **ETW Keyword Filter Property** dialog. From there, you can place a check mark in one or more of the **Keyword** check boxes for the selected ETW Provider, to cause the provider to return only the events that the **Keyword** enables. Whenever you enable a **Keyword**, you will see that the hexadecimal value in the text box at the bottom of the **ETW Keyword Filter Property** dialog changes to a new value. The default value is the 16-digit hexadecimal number: 0x0000000000000000, which typically signifies that all events for which the provider is configured will be delivered to Message Analyzer, that is, if the events are triggered and written to the ETW Session under which your provider is running during a trace.

More Information

To learn more about the differences between the **Keywords (Any)** and **Keywords (All)** settings, see Filtering with System ETW Provider Event Keywords and Levels.

Correlating Events Defined by Keywords

The **Keywords** that you enable will cause the ETW Provider to return only the events that the **Keywords** define. **Keyword** definitions and descriptions are usually specified in the manifest of the ETW Provider with which you are working. To be successful at specifying event **Keyword** configurations, you will need to understand the correlation between **Keywords** and the types of events that will be returned. You may be able to discover this correlative information by employing some of the methods described in Finding System ETW Provider Metadata. If you do not specify any **Keywords** for an ETW Provider, the provider will deliver all events that it is configured to provide, just as some ETW Providers that have no **Keyword** configurations do.

Support for WPP Tracing

This section describes how Message Analyzer provides support for processing WPP-generated events, in a Data Retrieval Session or a Live Trace Session.

Support for Parsing Static Trace Files

Message Analyzer can parse and display events that are generated by a Windows software trace preprocessor (WPP) trace provider. Because this type of provider writes events that can integrate with the ETW framework, Message Analyzer can load them from a saved event trace log (ETL) file that is created by an appropriate system

tool such as *Logman* or *Netsh*. To enable parsing of WPP-generated events, users must specify the path to a trace message format (TMF) file on the **WPP** tab of the **Options** dialog. If you have only a program database (PDB) file that provides the formatting information for the WPP event structure definitions, you can create a TMF file from it by specifying the path to the PDB file and the path to the utility *Tracepdb* on the **WPP** tab of the **Options** dialog, as described in Loading WPP-Generated Events.

Support for Parsing WPP Events Live

In scenarios where you want to view WPP events immediately after they are generated by a software component, it may be possible to capture WPP-generated events live through the use of an automatically-generated TMF file, if you specify a GUID for the WPP component that is generating the events. In this scenario, when the WPP component is compiled on a particular system, a TMF and/or a PDB symbol file will be automatically generated, for which Message Analyzer will need to search in order to locate the definition/s of the WPP event structure/s. This should enable Message Analyzer to parse the WPP events live.

To specify the GUID of the software component that is generating the events, you can enter it in the **Add Custom Provider** dialog, which is accessible from the **Add Providers** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog during Live Trace Session configuration.

More Information

To learn more about the configuration settings for system ETW Providers, including event **Keyword** and error **Level** filter configuration, see System ETW Provider Event Keyword/Level Settings.

To learn more about the ETW framework and system ETW Provider functionality, see the ETW Framework Conceptual Tutorial.

To learn more about how Message Analyzer supports WPP trace providers, see Loading WPP-Generated Events.

To learn more about Message Analyzer support for MOF-based providers, including how to register and deploy one, see Using MOF-Based ETW Providers.

Modifying Default Provider Settings

4 minutes to read

To enhance the capture configuration of a Live Trace Session, you can modify the following aspects of ETW-instrumented providers:

- **Driver/provider-level filtering** — as described in Common Provider Configuration Settings Summary, you can configure various low-level **Fast Filters** for PEF providers to introduce a selectivity factor to your Live Trace Sessions that improves performance, for example, with the **Microsoft-PEF-NDIS-PacketCapture** provider. You can also configure special stack filters and packet traversal paths at the driver level for **Trace Scenarios** that use the **Microsoft-Windows-NDIS-PacketCapture** provider. In addition, you can employ **WFP Layer Set** filters in **Trace Scenarios** that use the **Microsoft-PEF-WFP-MessageProvider** in order to control the direction in which packets are passed or blocked at the Transport Layer.

  Filtering at the driver/provider level improves speed because it reduces the number of messages that will be subject to the Message Analyzer Runtime parsing process, which includes **Session Filtering** that requires additional parsing — providing that a **Session Filter** is configured prior to running a Live Trace Session.

- **Event filtering** — as described in System ETW Provider Event Keyword/Level Settings, you can specify event **Keyword** and error **Level** filter settings for numerous system ETW Providers to enable selective event logging via ETW, which subsequently focuses the events that Message Analyzer captures.

Configuring Provider-Level Filters

The indicated provider setting types are accessible from the **Provider** or **ETW Core** tab of the **Advanced Settings** dialog that you can open by clicking the **Configure** link for message providers that display in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog in Live Trace Session configuration. The following is an overview of the filtering features that enable you to create the indicated configurations:

- **Fast Filter Groups** — configure **Groups** of logically chained **Fast Filters** and assign them to selected adapters in **Trace Scenarios** that use the **Microsoft-PEF-NDIS-PacketCapture** provider, for example, the **Local Network Interfaces** scenario on computers running the Windows 7 or Windows 8 operating system. Specify these settings in the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog during Live Trace Session configuration.

- **Multiple Fast Filters** — configure one or more **Fast Filters** in **Trace Scenarios** that use the **Microsoft-PEF-WFP-MessageProvider**, for example, the **Loopback and Unencrypted IPSEC** scenario. Specify these settings in the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog during Live Trace Session configuration.

- **WFP Layer Set Filters** — selectively enable or disable filters that pass or block inbound, outbound, or bidirectional packet traffic at the Transport Layer in **Trace Scenarios** that use the **Microsoft-PEF-WFP-MessageProvider**. Specify these settings in the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog during Live Trace Session configuration.

> **NOTE**
>
> You also have the option to log dropped packet information, which includes reason and layer statistics. Enable this function by placing a check mark in the **Select Discarded Packet Events** check box in the **Advanced Settings** dialog for the **Microsoft-PEF-WFP-MessageProvider**.

- **NDIS and Extension Stack Filters** — specify the layers on which packets are intercepted in the NDIS stack or Hyper-V-Switch extension stack, in **Trace Scenarios** that use the **Microsoft-Windows-NDIS-PacketCapture** provider, for example, the **Local Network Interfaces (Win8.1 and later)** scenario. You can also configure special filters such as **Truncation**, packet traversal **Direction**, **EtherTypes**, **IP Protocol Numbers**, **Mac Addresses**, and **IP Addresses**. Specify these filtering configurations in the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog during Live Trace Session configuration.
- **Hostname and Port Filters** — configure a **Hostname Filter** to restrict the capture of HTTP operations initiated by the client browser to requests and responses to and from a single web host, in **Trace Scenarios** that use the **Microsoft-PEF-WebProxy** provider, for example, the **Pre-Encryption for HTTPS** scenario. Configure a **Port Filter** to limit the traffic you capture to a particular HTTP port, such as 80, 443, or 8080. Specify these filtering configurations in the **Advanced Settings - Microsoft-PEF-WebProxy** dialog during Live Trace Session configuration.
- **ETW Event and Error Filters** — set the event **Keyword** and/or error **Level** configuration of system ETW providers by specifying settings on the **ETW Core** tab of the **Advanced Settings** dialog for any provider that displays in the **ETW Providers** list of the **New Session** dialog.

> **NOTE**
>
> For *system* ETW Providers (Windows components that have been instrumented as ETW event providers), the only additional configuration that you can specify are modifications to the event **Keyword** and error **Level** filtering settings; however, not all system ETW Providers make **Keyword** settings available for filtering. System ETW Providers, such as the **Microsoft-Windows-LDAP-Client**, are listed in the **Add System Providers** dialog, which is accessible from the **Add Providers** drop-down list on the **ETW Providers** toolbar on the **Live Trace** tab of the **New Session** dialog, as described in Adding a System ETW Provider.

More Information

To learn more about how to configure common provider settings, see the Common Provider Configuration Settings Summary.

To learn more about how to configure settings for the **Microsoft-PEF-NDIS-PacketCapture** provider, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

To learn more about how to configure settings for the **Microsoft-Windows-NDIS-PacketCapture** provider, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

To learn more about how to configure settings for the **Microsoft-PEF-WFP-MessageProvider**, see Using the Advanced Settings- Microsoft-PEF-WFP-MessageProvider Dialog.

To learn more about configuring ETW event **Keyword** and error **Level** settings, see System ETW Provider Event Keyword/Level Settings.

Common Provider Configuration Settings Summary

18 minutes to read

This topic provides an overview of the types of filters that you can apply to the core providers that are used in the most common **Trace Scenarios** that Message Analyzer provides. Examples of settings for such filters are included in the Common Provider Configuration Settings table of this topic.

Selecting Data with Trace Provider Filtering

Message Analyzer enables you to select specific data from a Live Trace Session by modifying the settings of ETW providers that are included in Message Analyzer **Trace Scenarios**. These settings enable you to configure key filter configurations for Message Analyzer providers in the following common **Trace Scenarios**:

- **Local Network Interfaces (Windows 8 and earlier)** — you can configure the following types of filtering for the **Microsoft-PEF-NDIS-PacketCapture** provider that is used in this scenario:
  - **Fast filtering** — enables you to filter message traffic based on offset-length patterns (OLPs), LinkLevelAddress, IPv4 address, and IPv6 address.
  - **Adapter filtering** — enables you to isolate traffic on specific adapters.
  - **Chained filtering** — enables you to create **Groups** of logically chained **Fast Filters** and assign them to specific adapters on your system.
- **Remote Network Interfaces** — you can configure the following types of filtering for the **Microsoft-Windows-NDIS-PacketCapture** provider:
  - **Adapter filtering** — enables you to isolate traffic to specific adapters, including remote host adapters and adapters for virtual machines (VMs) that are serviced by a Hyper-V-Switch.
  - **Driver filtering** — enables you to isolate inbound or outbound traffic to remote host adapters and to specify the NDIS filter layers on which to intercept packets that traverse the NDIS stack contained in the adapter.
  - **Switch filtering** — enables you to specify the filter layers on which to intercept packets that traverse a remote Hyper-V-Switch Extension stack and you can also specify the direction in which packets traverse this stack.
  - **Packet filtering** — other filters that you can specify to modify the trace configuration on a remote host with the **Microsoft-Windows-NDIS-PacketCapture** provider include **Truncation**, **EtherTypes**, **IP Protocol Numbers**, **MAC Addresses**, and **IP Addresses**.

Note: You can specify these same remote filtering configurations in *local* **Trace Scenarios** that use the **Microsoft-Windows-NDIS-PacketCapture** provider, for example, the **Local Network Interfaces (Windows 8.1 and later)** scenario. However, in this scenario, you will be applying the indicated filters to host or VM adapters that exist on the local computer by default.

You can review the user interface for setting these filtering configurations in the figure that is provided in Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

- **Loopback and Unencrypted IPSEC** — you can configure the following types of filtering for the **Microsoft-PEF-WFP-MessageProvider** that is used in this scenario:

  Note: The **Network Tunnel Traffic and Unencrypted IPSEC**, **Local Loopback Network**, and **SMB**

Client and Firewall scenarios also use the **Microsoft-PEF-WFP-MessageProvider** and can utilize the same type of filtering specified below.

- **Transport Layer filtering** — enables you to isolate inbound and outbound Transport Layer message traffic for IPv4 and IPv6 transports, by using the **WFP Layer Set** filters.
- **Fast filtering** — enables you to filter messaging traffic based on IPv4 address, IPv6 address, TCP port, and UDP port values.

Note: You can also choose to select discarded packet events for logging.

You can review the user interface for setting these filtering configurations in the figure that is provided in Using the Advanced Settings- Microsoft-PEF-WFP-MessageProvider Dialog.

- **Pre-Encryption for HTTPS**— you can configure the following types of filtering for the **Microsoft-PEF-WebProxy** provider that is used in this scenario:
  - **HTTP filtering** — enables you to specify a **Hostname Filter** and/or a **Port Filter** so that you can isolate HTTP message data to specified host names and port numbers.

Note: You can also specify a certificate file to use for server authentication on HTTPS connections, as required by some sites. In addition, you can keep client or server connections alive for performance improvements, by selecting the **Reuse Client Connections** and/or **Reuse Server Connections** check boxes on the **Provider** tab of the **Advanced Settings - Microsoft-PEF-WebProxy** dialog, which is accessible in the previously described manner.

Fast Filters

You can modify the provider configuration in the **Local Network Interfaces (Win 8 and earlier)**, **Local Loopback Network**, and **Loopback and Unencrypted IPSEC Trace Scenarios** by specifying low-level **Fast Filters**. Message Analyzer enables you to add simple filtering values as part of your **Fast Filter** configuration, as described in the Common Provider Configuration Settings table. **Fast Filters** reduce the data volume that is transported from kernel to user mode (Runtime parsing engine), which saves system resources and also improves Message Analyzer efficiency under heavy traffic loads.

You can apply **Fast Filters** in any **Trace Scenario** that uses the **Microsoft-PEF-WFP-MessageProvider** or the **Microsoft-PEF-NDIS-PacketCapture** provider. **Fast Filters** for the **Microsoft-PEF-WFP-MessageProvider** are configurable from the **Provider** tab of the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog that is accessible by clicking the **Configure** link for the **Microsoft-PEF-WFP-MessageProvider** that displays in the **ETW Providers** list when you select one of the applicable **Trace Scenarios** in the **New Session** dialog. **Fast Filters** for the **Microsoft-PEF-NDIS-PacketCapture** provider are configurable from the **Provider** tab of the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog that is accessible by clicking the **Configure** link for the **Microsoft-PEF-NDIS-PacketCapture** provider that displays in the **ETW Providers** list when you select a **Local Network Interfaces Trace Scenario** in the **New Session** dialog. These filters enable you to focus a provider's message retrieval action at a lower level than a **Session Filter**, which makes them faster and more efficient.

For example, you can configure up to three **Fast Filters** per **Group** in the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, and each filter can be any of the four filter types described earlier in Selecting Data with Trace Provider Filtering. If you specify more than one **Fast Filter** per **Group**, then you can choose to logically AND or OR the filters together. Thereafter, you can assign the filter **Groups** to specific adapters that are listed in the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog. Also note that filter **Groups** are logically ANDed together by default.

You can learn more about configuring **Groups** of **Fast Filters** and adapter assignment in **Trace Scenarios** that use the **Microsoft-PEF-NDIS-PacketCapture** provider by reviewing the topic Using the Advanced Settings -

Microsoft-PEF-NDIS-PacketCapture Dialog. However, some of the settings for adapter and **Fast Filter** configuration are briefly described in the configuration settings table in this section.

Applying Filters in Local and Remote Scenarios

This section provides a quick overview of applying host adapter filtering and low-level **Fast Filters** in **Trace Scenarios** that capture traffic locally; and host adapter filtering, packet filters, and other special filters in **Trace Scenarios** that capture traffic remotely. To specify these filtering configurations, you will need to click the **Configure** link in the **ETW Providers** list in the **New Session** dialog for the particular provider with which you are working in order to display the associated **Advanced Settings** dialog.

Local Trace Scenarios

You can modify the **Microsoft-PEF-NDIS-PacketCapture** provider configuration for the **Local Network Interfaces** scenario in the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, to isolate traffic by the following:

- Selecting a particular local adapter on which to capture traffic.
- Specifying the direction in which traffic is captured, either incoming, outgoing, or in both directions.
- Configuring up to two **Fast Filter Groups**, each with up to three **Fast Filters** specified.

Local Filtering Configuration — you can specify filtering configuration settings in the **Advanced Settings** dialog for the **Microsoft-PEF-NDIS-PacketCapture** provider whenever you run the **Local Network Interfaces** scenario to capture data on the local computers running the Windows 7 or Windows 8 operating system.

For example, you can effectively create an adapter filter by selecting at least one or both of the **In** and **Out** check boxes for a chosen local host adapter. in the **System Network** section of the dialog. This action will define the adapter that passes the packet traffic and in the direction/s that you specified.

You can also create one or two logically-related **Fast Filter Groups** and apply the configuration to a selected adapter to isolate specific data based on the filtering criteria, as described in Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

Remote Trace Scenarios

You can specify filtering configuration settings in the **Advanced Settings** dialog for the **Microsoft-Windows-NDIS-PacketCapture** provider whenever you run the **Remote Network Interfaces** or **Remote Network Interfaces with Drop Information Trace Scenarios**, to isolate traffic by the following:

- Selecting remote host adapters on which to capture traffic.
- Selecting a remote virtual machine (VM) adapter that is serviced by a Hyper-V-Switch, on which to capture traffic.
- Specifying the traffic direction and NDIS filter layers on which to intercept packet traffic for remote host adapters.
- Specifying the packet traversal path and Hyper-V-Switch extension layers on which to intercept packet traffic.
- Specifying other special filters that modify the scope of data retrieval, such as **Truncation**, **EtherType**, **IP Protocol Number**, **MAC Address**, and **IP Address** filters.

Tip: You can also apply these settings in other **Trace Scenarios** that use the **Microsoft-Windows-NDIS-PacketCapture** provider, which includes capturing data on a local host.

Remote Filtering Configuration — you can specify filtering configuration settings in the **Advanced Settings** dialog for the **Microsoft-Windows-NDIS-PacketCapture** provider, after you select any **Trace Scenario** that

uses the **Microsoft-Windows-NDIS-PacketCapture** provider to capture data on a target remote (or local) computer.

For a detailed explanation of all the filtering configurations that you can apply for the **Microsoft-Windows-NDIS-PacketCapture** provider, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

Keyword Filters — although you can specify event **Keyword** and error **Level** filters for any provider that defines them, you have the option to specify a considerable number of **Keyword** filters from the **ETW Core** tab of the **Advanced Settings** dialog for the **Microsoft-Windows-NDIS-PacketCapture** provider in both local and remote scenarios. For example, you can specify event **Keywords** for dropped packets, response time, and diagnostics.

For further details on applying event **Keyword** and error **Level** filters, see System ETW Provider Event Keyword/Level Settings.

More Information
To learn more about capturing data from a remote host, see the Configuring a Remote Capture section. However, some of the configuration settings for remote tracing are briefly described in the configuration settings table that is included in this section.

Common Provider Configuration Settings

The table that follows provides the details for specifying filtering configuration settings for ETW message providers in the following common **Trace Scenario** types:

- **Local Network Interfaces** — using the **Microsoft-PEF-NDIS-PacketCapture** provider.
- **Remote Network Interfaces** — using the **Microsoft-Windows-NDIS-PacketCapture** provider.
- **Loopback and Unencrypted IPSEC** — using the **Microsoft-PEF-WFP-MessageProvider**.
- **Pre-Encryption for HTTPS** — using the **Microsoft-PEF-WebProxy** with Fiddler provider.

Note: The built-in **Trace Scenarios** included with Message Analyzer can contain a combination of PEF and Windows system ETW providers. By understanding how to configure settings for these provider types, as used in the common **Trace Scenarios** and described in the table that follows, you will also understand how to modify the providers in other default **Trace Scenarios**, as appropriate to your requirements.

Table 5. Common Provider Configuration Settings

DEFAULT TRACE SCENARIOS	CONFIGURATION SETTINGS	PROPERTY OR FEATURE	DESCRIPTION

DEFAULT TRACE SCENARIOS	CONFIGURATION SETTINGS	PROPERTY OR FEATURE	DESCRIPTION
Local Network Interfaces Capture local Link Layer traffic from NDIS	**Adapters**	**In, Out**	This feature is accessible from the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog. By selectively enabling or disabling the **In** and **Out** check boxes for adapters listed in this dialog, you can both enable specific adapters through which to capture message traffic and the direction in which to capture it. The adapters you can select in the dialog are enumerated on your system when the **Microsoft-PEF-NDIS-PacketCapture** provider is installed. By default, Message Analyzer enables local adapters for capturing data in both directions. To prevent an adapter from capturing message traffic, simply deselect its **In** and **Out** check boxes. **Note:** You can also isolate message traffic to a specific adapter by configuring a **Fast Filter** to contain its **LinkLevelAddress**, as described in the Fast Filters section of this table.
Local Network Interfaces Capture local Link Layer traffic from NDIS	**Fast Filters**	**Filter type**	This feature is accessible from the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog. It enables you to use a drop-down list to select the type of filters you want to configure, a **Group** to contain them, and the adapters to assign them to when configuring a **Local Network Interfaces** scenario. The **Fast Filters** that you can specify consist of the following: - **OLP** — enables you to specify an offset length pattern (OLP) that filters for messages containing the pattern. For example, you might set an OLP to *34:8:7F* to return only messages that contain this particular OLP.

DEFAULT TRACE SCENARIOS	CONFIGURATION SETTINGS	PROPERTY OR FEATURE	DESCRIPTION
			Operators that function with the OLP filter type consist of equal (=), not equal (!=), less than (<), and greater than (>). **Note:** As a result of complexities and the impact on performance, you should only configure a single OLP filter per adapter in the **Microsoft-PEF-NDIS-PacketCapture** provider settings, although different adapters can have different filters. **More Information** **To learn more** about OLP filtering, see OLP Filters. - **LinkLevelAddress** — enables you to specify a Media Access Control (MAC) hardware address for which inbound and outbound traffic is targeted to a particular adapter device. You can specify a **LinkLevelAddress** in a six-byte, hexadecimal format, similar to the following example: *01-23-45-67-89-AB*. - **IPv4Address** — enables you to isolate message traffic to a specified IPv4 address, by specifying a value such as the following in the text box to the right of the **Filter** type drop-down menu: *192.168.1.1* - **IPv6Address** — enables you to isolate message traffic to a specified IPv6 address, by specifying a value such as the following in the text box to the right of the **Filter** type drop-down menu: *fe80::9de5:fc31:8856:58a8%11* **More Information** **To learn more** about configuring **Fast Filters** for the **Microsoft-PEF-NDIS-PacketCapture** provider and assigning them to specific adapters, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

DEFAULT TRACE SCENARIOS	CONFIGURATION SETTINGS	PROPERTY OR FEATURE	DESCRIPTION
Remote Network Interfaces Remote capture on Link Layer	**Adapters**	**Enabling**	This feature is accessible from the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog. By selecting the **Enabled** check boxes in this dialog for host adapters and/or adapters for virtual machines (VMs) that are serviced by a Hyper-V-Switch, you can specify the adapters through which to capture remote traffic, and in so doing, enable those adapters to receive any filtering configurations that you specify in the **Advanced Settings** dialog. **Important:** If you want to capture traffic from a specific remote VM, you will need to select the VM in the Interface Selection section of the **Advanced Settings** dialog and then specify a filter based on the VM's **MAC Address** to isolate the data. Otherwise, you will capture Hyper-V-Switch traffic that is destined for all VMs that are serviced by the switch, given that a Hyper-V-Switch driver cannot of itself distinguish between VMs. The adapters that you can select in the dialog are enumerated on your system when you connect to a specified remote host. By default, Message Analyzer enables all adapters in the **Advanced Settings** dialog. To prevent traffic from being captured on a listed adapter, you can simply deselect it. **Note:** You can also capture remote (or local) traffic in Promiscuous Mode, for adapters that support it, by simply selecting a **P-Mode** check box for a supporting adapter.

DEFAULT TRACE SCENARIOS	CONFIGURATION SETTINGS	PROPERTY OR FEATURE	DESCRIPTION
Remote Network Interfaces Remote capture on Link Layer	**Filters**	The types of filtering configurations that you can apply in remote tracing scenarios consist of the following: - **All Layers** filter - Packet **Truncation** filter - **Direction** filter – - **Ingress/Egress** packet direction path through NDIS filter stack when applied to host adapters. - **Ingress/Egress** packet traversal path on the Hyper-V-Switch extension stack when applied to a switch adapter. - **EtherType** filter - **IP protocol number** filter - **MAC Address** filter - **IP Address** filter	The configuration for these filters is located in the **Filters** pane of the **Advanced Settings** dialog. From the dialog, you can choose the NDIS filter layers on which packets are intercepted in remote host adapters or on the Extension layers of a Hyper-V-Switch that services a remote VM adapter, for troubleshooting purposes. You can also set the direction in which to capture data on remote host adapters, and the packet traversal paths through a remote Hyper-V-Switch Extension stack. Other filters that you can specify include **Truncation**, **EtherType**, **IP Protocol Numbers**, **MAC Address**, and **IP Address** filters. **More Information** **To learn more** about how to configure such filters for a remote trace, see the Configuring a Remote Capture and Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog topics.

DEFAULT TRACE SCENARIOS	CONFIGURATION SETTINGS	PROPERTY OR FEATURE	DESCRIPTION
Loopback and Unencrypted IPSEC Windows Filtering Platform Tracing	**WFP Layer Set**		Provides the following **WFP Layer Set** filters for **Trace Scenarios** that use the **Microsoft-PEF-WFP-MessageProvider**, so that you can control the direction in which packets are configured to pass at the Transport Layer for IPv4 and IPv6 traffic: - **INBOUND_TRANSPORT_V4** - **OUTBOUND_TRANSPORT_V4** - **INBOUND_TRANSPORT_V6** - **OUTBOUND_TRANSPORT_V6**
Loopback and Unencrypted IPSEC Windows Filtering Platform Tracing	**INBOUND_TRANSPORT_V4**	**Check box select/unselect**	A kernel-mode TCP/IP stack filter that operates in the *receive* path at the Transport Layer before any processing occurs at that layer. This layer is above IPsec processing at the Network layer and below Application Level Enforcement (ALE) layers. Therefore when selected, this filter enables capture of all *inbound* packets at the Transport Layer, with the exclusion of any that are dropped at the Network layer. You can select or unselect this filter to capture or not capture TCP/IPv4 packets, respectively, at the Transport Layer.

DEFAULT TRACE SCENARIOS	CONFIGURATION SETTINGS	PROPERTY OR FEATURE	DESCRIPTION
Loopback and Unencrypted IPSEC Windows Filtering Platform Tracing	**OUTBOUND_TRANSPORT_V4**	**Check box select/unselect**	A kernel-mode TCP/IP stack filter that operates in the *send* path at the Transport Layer before any processing occurs at that layer. This layer is above IPsec processing at the Network layer and below Application Level Enforcement (ALE) layers. Therefore when selected, this filter enables capture of all *outbound* packets at the Transport Layer, with the exclusion of any that are dropped at the Network layer. You can select or unselect this filter to capture or not capture TCP/IPv4 packets, respectively, at the Transport Layer.
Loopback and Unencrypted IPSEC Windows Filtering Platform Tracing	**INBOUND_TRANSPORT_V6**	**Check box select/unselect**	Functionally similar to the inbound IPv4 version of this filter type, this is a kernel-mode TCP/IP stack filter layer in the *receive* path that you can select or unselect to capture or not capture inbound TCP/IPv6 packets, respectively, at the Transport Layer.
Loopback and Unencrypted IPSEC Windows Filtering Platform Tracing	**OUTBOUND_TRANSPORT_V6**	**Check box select/unselect**	Functionally similar to the outbound IPv4 version of this filter type, this is a kernel-mode TCP/IP stack filter layer in the *send* path that you can select or unselect to capture or not capture inbound TCP/IPv6 packets, respectively, at the Transport Layer.

DEFAULT TRACE SCENARIOS	CONFIGURATION SETTINGS	PROPERTY OR FEATURE	DESCRIPTION
Loopback and Unencrypted IPSEC Windows Filtering Platform Tracing	**Fast Filters** [2]	**FilterType**	Improves performance — for example, when processing large volumes of traffic — by enabling you to filter out unwanted traffic via IP address and port filters. You can specify the type of **Fast Filter** you want to apply to a **Loopback and Unencrypted IPSEC** trace, by selecting it from a drop-down list that includes the following items: - **IPv4** —should be specified in a format similar to the following: *192.168.1.1*. - **IPv6** — should be specified in a format similar to the following: *2001:4898:2b:3:d824:99e9:7371:31d9* or *fe80::6e9c:edff:fe94:ec0011* - **TCP port** — should be specified in integer format, for example: *80*. - **UDP port** — should be specified in integer format, for example: *53*. **Note:** For tracing with the **Microsoft-PEF-WFP-MessageProvider**, the specified filter is applied only to the corresponding layer(s). For example: *IPv4Address==192.168.1.1* is applied to Transport_V4 inbound and outbound layers. There is no filter on the V6 counter parts. If the V6 layers are enabled, you will see frames of IPv4 and IPv6 packets. If you want to see only IPv4 messages, then enable only the IPv4 layers in the **WFP Layer Set**; then configure a **Fast Filter** with an **IPv4** address. TCP and UDP port numbers should be integers between 0 and 65535.

DEFAULT TRACE SCENARIOS	CONFIGURATION SETTINGS	PROPERTY OR FEATURE	DESCRIPTION
Loopback and Unencrypted IPSEC Windows Filtering Platform Tracing		**Filter text box**	Provides the entry point where you specify a value for the type of filter that you selected in a **Fast Filter** drop-down list.
Pre-Encryption for HTTPS Capture HTTPS client-side unencrypted traffic with the **Microsoft-Pef-WebProxy** with Fiddler provider	**HostnameFilter**	Host address in the format: *www.bing.com*	Filters HTTP packets from a web server based on the host name.
Pre-Encryption for HTTPS Capture HTTPS client-side unencrypted traffic with the **Microsoft-Pef-WebProxy** with Fiddler provider	**PortFilter**	Port number in a format similar to the following: *80*	Filters packets by numbered ports only.
Pre-Encryption for HTTPS Capture HTTPS client-side unencrypted traffic with the **Microsoft-Pef-WebProxy** with Fiddler provider	**HTTPS Client Certificate**	A certificate file in *.cer format.	Specifies the .cer file for Fiddler to provide to a web server that requires certification validation.
Pre-Encryption for HTTPS Capture HTTPS client-side unencrypted traffic with the **Microsoft-Pef-WebProxy** with Fiddler provider	**Reuse Server Connections** **Reuse Client Connections**	**Enabling check boxes**	When selected, keeps server and/or client connections alive to improve performance, rather than re-creating new connections.
System ETW Providers	Event **Keyword** filter	**Keywords(All)** **Keywords(Any)**	Select event **Keywords** from the **ETW Keyword Filter Property** dialog that is accessible from the **ETW Core** tab of the **Advanced Settings** dialog for a particular message provider. Predefined **Keywords** in the **ETW Keyword Filter Property** dialog are selectable by **Value** name and translate to 16-digit hexadecimal numbers, for example: *0x00000000000000C0*.

DEFAULT TRACE SCENARIOS	CONFIGURATION SETTINGS	PROPERTY OR FEATURE	DESCRIPTION
System ETW Providers	**Level** filter	**LogAlways** **Critical** **Error** **Warning** **Information** **Verbose**	Settings enable filtering based on the severity or verbosity of error events., or in the case of the **LogAlways** level, ensures that logging of all **Levels** will always occur.

[2] When capturing data in **Trace Scenarios** that use the **Microsoft-PEF-WFP-MessageProvider**, setting a **Fast Filter** such as **IPAddress** == *192.168.1.1* does not prevent IPv6 traffic from being captured because IPv4 and IPv6 messages are retrieved from separate layers in these scenarios and the **Fast Filter** applies only to the IPv4 layer.

More Information

To learn more about creating **Fast Filter** and adapter filter configurations for a Live Trace Session that uses the **Local Network Interfaces Trace Scenario**, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

To learn more about how to select adapters and specify filters for a Live Trace Session that uses the **Remote Network Interfaces Trace Scenario**, see Configuring a Remote Capture and Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

To learn more about configuring system ETW Providers with event **Keyword** and **Level** filters, see System ETW Provider Event Keyword/Level Settings.

See Also

PEF Message Providers

Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog

9 minutes to read

The **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog is available in Message Analyzer installations on computers that are running the Windows 7 or Windows 8 operating system only. The dialog is accessible by clicking the **Configure** link to the right of the **Microsoft-PEF-NDIS-PacketCapture** provider **Id** in the **ETW Provider** list on the **Live Trace** tab of the **New Session** dialog; that is, after you select a **Trace Scenario** that uses the **Microsoft-PEF-NDIS-PacketCapture** provider. For example, you can select any of the following **Trace Scenarios** from the **Select Scenarios** drop-down list to display the **Microsoft-PEF-NDIS-PacketCapture** provider in the **ETW Providers** list of the **New Session** dialog during Live Trace Session configuration:

- **Local Network Interfaces (Win 8 and earlier)**
- **Wired Local Area Network (Win 8 and earlier)**
- **Wireless Local Area Network (Win 8 and earlier)**

The **Advanced Settings** dialog for the **Microsoft-PEF-NDIS-PacketCapture** provider is shown in the figure that follows.

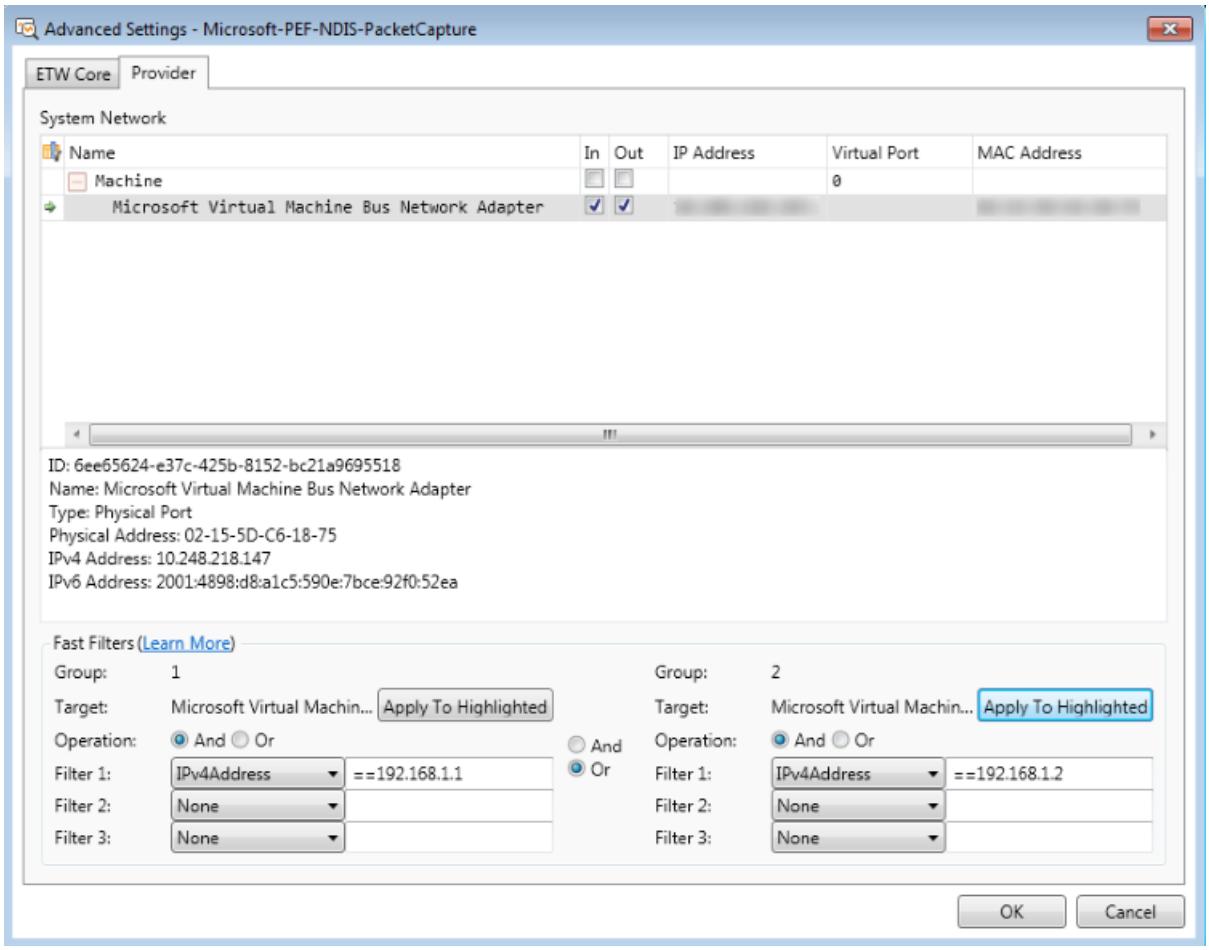

Figure 22: Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog

The **Advanced Settings** dialog provides its functionality in a flexible framework that enables you to focus on capturing very specific data while achieving the performance advantages that are inherent to **Fast Filters**.

This framework enables you to logically chain up to three **Fast Filters** in each of two **Groups** that can be logically ORed or ANDed, which you can then assign to one or more *selected* adapters. For example, you can assign a single **Group** or both **Groups** of filters to one or more adapters.

An adapter is considered selected only when you select at least one of the **In** or **Out** (traffic direction) check boxes in the **Advanced Settings** dialog. However, the default selection is to capture traffic in both directions. In order for a **Group** filter configuration to *apply* to packets intended for a particular adapter, the adapter must be selected in the indicated manner. Also, a filter **Group** can only be *assigned* to a selected adapter if the row in which the adapter exists is highlighted and then you click the **Apply to Highlighted** button for the **Group**. In addition, because **Fast Filters** are nested in **Groups**, which are in turn assigned to specific adapters, the total filtering effects that you realize are the result of both adapter selection and filter **Group** assignment combined. For this reason, you should carefully consider how you assign **Groups** to selected adapters, as described in How Fast Filter Groups Are Applied.

Viewing the System Network Configuration

When you open the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, the **System Network** tree grid configuration is prepopulated with the following information:

- **Adapters** — the adapters and adapter nodes on your machine are listed in the **Name** column of the tree grid configuration. For example, your adapters might include Ethernet and Wireless network adapters and a Hyper-V-Switch node containing virtual machines (VMs).
- **Traffic direction** — for each adapter on your system there are **In** and **Out** check boxes that enable you to specify the direction of the traffic you want to capture, for example, inbound to a particular adapter or outbound from the adapter. All the check boxes are selected by default to enable you to capture data in both directions, which is the typical configuration for the most useful context. However, there can be times when you want to isolate traffic in a particular direction. In this case, you can select the **In** and **Out** check boxes individually for any listed adapter. Otherwise, you can globally select inbound and outbound traffic for all listed adapters by selecting the **In** and **Out** check boxes in the **Machine** or **Adapters** row of the **System Network** tree grid configuration.
- **IP addresses** — the IP address of each adapter on your machine is specified in the **IP Address** column of the tree grid configuration, in both IPv4 and IPv6 formats.
- **Virtual ports** — any applicable ports for local VM adapters are listed in the **Virtual Port** column of the tree grid configuration.
- **MAC addresses** — the MAC address for each adapter on your machine is listed in the **MAC Address** column of the tree grid configuration.

Finding Column Data

If you have a particularly long list of adapters, including a Hyper-V-Switch and VMs, you can take advantage of the **Column Filter** feature that is also included in numerous Message Analyzer viewer and **Tool Windows**, to search for column entries that contain specified search text. This can help you to quickly isolate a row of adapter data in which you are interested. You can display the column filtering row by clicking the **Show Column Filter Row** icon in the **System Network** configuration. Also, if you select any network adapter that is listed, a description displays below the **System Network** configuration.

> **NOTE**
>
> You can utilize the **Copy** commands in the context menu that displays when you right-click any tree grid configuration column that contains data.

Logically Chaining Fast Filters

You have the option to configure up to three **Fast Filters** per **Group** with the same or different address types, and you can also logically chain the **Fast Filters** within each **Group**, by selecting either the **And** or **Or** operator option. Note that the **Groups** are logically ANDed by default, but can be set to **OR** if necessary.

Assigning Fast Filter Groups to Tree Grid Elements

The **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog provides a flexibility that enables you to selectively assign filter **Groups** to adapters in the **System Network** tree grid. After you create a **Fast Filter** configuration for a particular **Group**, you can assign it to any of the following elements in the tree grid by first highlighting the element and then clicking the appropriate **Apply to Highlighted** button:

- **One adapter element row** — when you assign a filter **Group** containing one or more **Fast Filters** to a single adapter that is both selected and highlighted, the filtering configuration applies to the highlighted adapter only.

> **NOTE**
> You can highlight an adapter by clicking any column in the element row in which the target adapter is listed, at which time the row is highlighted in blue. After you assign a filter **Group** to a highlighted adapter by clicking the **Apply to Highlighted** button, the color of the highlighted row changes to gray.

- **Multiple adapter element rows** — when you assign a filter **Group** of one or more **Fast Filters** to multiple selected and highlighted adapters, the filtering configuration applies to the highlighted adapters.
- **All adapters** — when you assign a filter **Group** of one or more **Fast Filters** to all selected and highlighted adapters, the filtering configuration applies to all adapters in the **System Network** tree grid.

> **NOTE**
> You can select and highlight all adapters by selecting one or both traffic direction check boxes in the **Machine** element row.

Also, when you successfully assign a filter **Group** to a particular adapter, the name of the adapter to which the filter configuration is assigned displays next to the **Target** label for the assigned **Group**.

How Fast Filter Groups Are Applied

In the grouping model, filter **Group** behavior varies depending on how you assign the **Groups** and to some extent on the logical operator with which the **Groups** are chained. For a simple case, when you click the **Apply to Highlighted** button/s to assign one or both **Groups** to a single adapter or a top node containing child adapters in the **System Network** tree grid, a property is set that designates the adapter/s as the target for the **Group/s**. Thereafter, any packet that is intended for the specific adapter/s will be retrieved by the **Microsoft-PEF-NDIS-PacketCapture** provider if it passes the filtering criteria contained in the assigned **Groups**. Packets that are not intended for the specific adapter/s are not affected by the assigned filter **Groups**. Moreover, if you assign each **Group** to a *different* adapter or node, any arriving packet that applies to both nodes or adapters causes the rules of both **Groups** to fire in accordance with the defined operators. Under these circumstances, if a packet arrives that applies to only one node or adapter, then the rules fire for only the **Group** assigned to that node or adapter. In this case, the logical AND that chains the **Groups** by default has no effect.

The main scenario for which filter **Groups** were developed in the **Advanced Settings** dialog has to do with local VMs that are serviced by a Hyper-V-Switch and the difficulty of isolating messages to a particular VM. Because Message Analyzer enumerates Hyper-V networks, you can now create a **Group** filter configuration and assign it to a specific VM, rather than having to create a **Fast Filter** for a specific VM and then capture traffic locally on that VM.

In the grouping model, if a packet arrives that belongs to a particular element such as **Machine**, or a Hyper-V-Switch, VM, or physical adapter, and a filter **Group** is assigned to that element, then the filtering rules configured in the assigned **Group** are applied. However, packets can belong to more than one element. For instance, packets always belong to the **Machine** element, but might also belong to a particular Hyper-V-Switch or a VM that is serviced by the switch. When this is the case, the grouping model creates a flexibility that enables you to isolate packets to whatever element you choose, including individual VMs. For example, you could have the following scenarios and interactions:

- **Group 1** is assigned to a VM; **Group 2** is assigned to a physical switch (with no Hyper-V network) — every packet that arrives either belongs to the VM, physical switch, or some other element, but never to both. Therefore, the default logical AND that is applied to the **Groups** has no effect because both **Groups** can never apply to the same packet. In this situation, when a packet arrives that belongs to the VM **Group**, the filtering rules for **Group 1** fire; and when a packet arrives that belongs to the physical switch **Group**, the filtering rules for **Group 2** fire.
- **Group 1** and **Group 2** are assigned to the same element/node — under these conditions, the default logical AND that is assigned to the **Groups** simply defines the top-level operator and subsequently the behavior of the Boolean expression you are constructing. When a packet arrives that applies to the particular element to which the **Groups** are assigned, the filtering rules for both **Groups** fire.
- **Group1** is assigned to **Machine**; **Group2** is assigned to a VM — in this case, if a packet arrives that belongs to the VM, then in effect it belongs to both **Groups** (all packets belong to **Machine**), and therefore the filtering rules for both **Groups** kick in with the defined operators. If a packet arrives that does not belong to the VM, then only the filtering rules for **Group 1** fire.

More Information
To learn more about how to create a **Fast Filter Group** and assign it to an adapter, see the procedure Configure and Run a Local Network Interfaces Trace.

See Also

Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog

Using the Advanced Settings- Microsoft-PEF-WFP-MessageProvider Dialog

4 minutes to read

This section describes the types of filters that you can configure for the **Microsoft-PEF-WFP-MessageProvider**. The filtering configurations are available from the **Advanced Settings** dialog for this provider, which displays after you click the **Configure** link to the right of the **Microsoft-PEF-WFP-MessageProvider** after it appears in the **ETW Providers** list of the **New Session** dialog during Live Trace Session configuration. The **New Session** dialog is accessible from the Message Analyzer **File** menu, the **Start Page**, or the global Message Analyzer toolbar.

You can add the **Microsoft-PEF-WFP-MessageProvider** to the **ETW Providers** list, by selecting any **Trace Scenario** that uses this provider from the **Select Scenario** drop-down list in the **New Session** dialog during Live Trace Session configuration. For example, you can select any of the following **Trace Scenarios** to display the **Microsoft-PEF-WFP-MessageProvider** in the **ETW Providers** list:

- **Loopback and Unencrypted IPSEC**
- **Local Loopback Network**
- **Network Tunnel Traffic and Unencrypted IPSEC**
- **SMB2 Client and Firewall**

The **Advanced Settings** dialog for the **Microsoft-PEF-WFP-MessageProvider** is shown in the figure that follows:

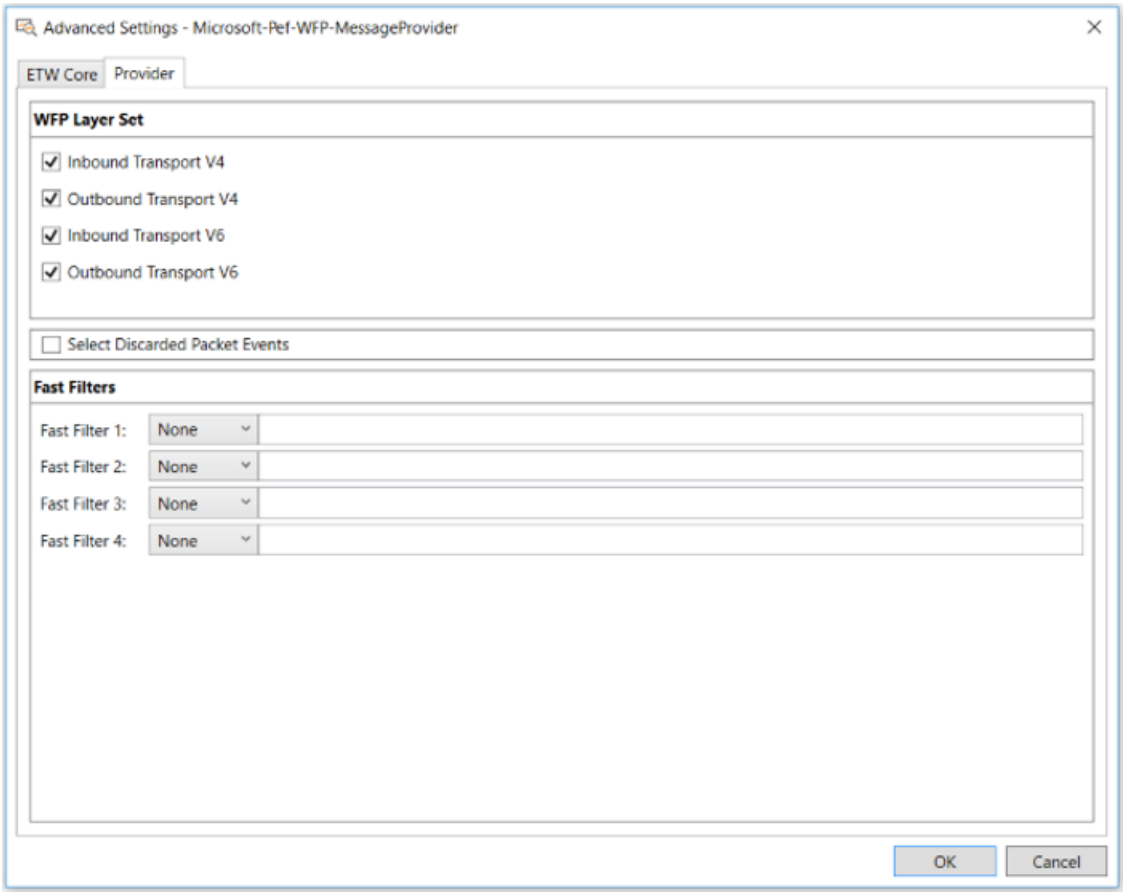

Figure 23: Advanced Settings for the Microsoft-PEF-WFP-MessageProvider

Using WFP Layer Set Filters

When configuring settings for the **Microsoft-PEF-WFP-MessageProvider** in **Trace Scenarios** that use it, you can create a **WFP Layer Set** filter configuration that directionally isolates inbound or outbound TCP packets at the Transport Layer for IPv4 or IPv6 traffic. A scenario where you might want to do this is when you are capturing loopback traffic, in which case, you should disable either inbound or outbound IPv4 and IPv6 traffic to avoid the duplication of messages. Otherwise, for regular network traffic, you should always enable both inbound and outbound packet directions.

You can specify the configuration for the **WFP Layer Set** filters on the **Provider** tab of the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog. The **WFP Layer Set** contains the following filters, which you can enable or disable, respectively, by selecting or unselecting the corresponding filter check boxes as appropriate:

- INBOUND_TRANSPORT_V4
- OUTBOUND_TRANSPORT_V4
- INBOUND_TRANSPORT_V6
- OUTBOUND_TRANSPORT_V6

These check boxes represent kernel mode TCP/IP stack filters that operate in the receive or send path (inbound or outbound, respectively) at the Transport Layer before any processing occurs at this layer. When you create a **WFP Layer Set** configuration with these check boxes, you selectively enable or disable the kernel mode filters that pass or block inbound, outbound, or bidirectional packet traffic at the Transport Layer, depending on the settings, for IPv4 and IPv6 traffic.

Using WFP Fast Filters

The actual work that is performed by the **Fast Filters** that you specify in the **Microsoft-PEF-WFP-MessageProvider** configuration is accomplished by the WFP base filtering engine (BFE). The message frames that pass the filtering criteria are delivered to the **Microsoft-PEF-WFP-MessageProvider** callout drivers at the corresponding layers, which in turn send the messages to the enabling ETW session.

WFP Fast Filters consist of the following types:

- **IPv4** — enables you to filter live traffic based on a specified IPv4 address, as described in IPv4Address Filters.
- **IPv6** — enables you to filter live traffic based on a specified IPv6 address, as described in IPv6Address Filters.
- **TCP port** — enables you to filter live traffic based on a specified TCP port.
- **UDP port** — enables you to filter live traffic based on a specified UDP port.

Utilizing Fast Filter Relational Operators

When you configure **Fast Filters** in the **Advanced Settings - Microsoft-PEF-WFP-MessageProvider** dialog, you have the option to enhance filter functionality with the use of several relational operators, which includes the following:

- Logical NOT (!)
- GREATER THAN (>)
- LESS THAN (<)

For example, you might apply **Fast Filter** configurations that exclude or include specific traffic on a **TCP port** or **UDP port**, respectively. To do this, you could configure one **Fast Filter** with a **TCP port** set to remove HTTPS

traffic from your live trace and another **Fast Filter** with a **UDP port** set to exclude all traffic except LDAP messages, respectively, as follows:

Fast Filter 1 value for **TCP port**: `!= 443`
Fast Filter 2 value for **UDP port**: `== 389`

Logging Dropped Packets

The **Microsoft-PEF-WFP-MessageProvider** also enables you to log dropped packet information, which includes reason and layer statistics. To obtain these statistics, you must select the **Select Discarded Packet Events** check box on the **Provider** tab of the **Advanced Settings – Microsoft-PEF-WFP-MessageProvider** dialog.

Logged information for discarded packet events consists of dropped packet statistics that are derived from the Discard filter layer of the Statistics callout in the **Microsoft-PEF-WFP-MessageProvider**. Statistics consist of the reason for dropping packets and the layer on which they were dropped.

You can view the dropped packet statistics as ETW events in the **Analysis Grid** viewer after a trace with the **Microsoft-PEF-WFP-MessageProvider** is complete. You can use the **Column Filter** feature for the **Summary** column of the **Analysis Grid** viewer to specify a search term such as "discard" to expose any messages that might indicate that the firewall was involved in blocking traffic. The **Discarded Packet Events** feature can help you troubleshoot whether packets are being dropped by the network, the **Microsoft-PEF-WFP-MessageProvider**, or the firewall.

> **NOTE**
> When you select the **Select Discarded Packet Events** check box on the **Provider** tab of the **Advanced Settings** dialog, any **Fast Filters** or **WFP Layer Set** filters that you have specified will not be applied to the discarded packet events.

More Information
To learn more about using **Column Filters**, see Filtering Column Data.

See Also

Microsoft-PEF-WFP-MessageProvider

System ETW Provider Event Keyword-Level Settings

14 minutes to read

All PEF providers are instrumented with Event Tracing for Windows (ETW) technology so that Message Analyzer can leverage its infrastructure for data collection, session control, buffer configuration, and so on, as described in the ETW Framework Conceptual Tutorial. As a result, all PEF providers contain a core ETW Provider component that interacts with an enabling ETW Session where it writes events that Message Analyzer can capture. Other ETW Providers that are registered on your system were originally created by instrumenting various Windows components with ETW technology; as a result, they too can leverage the ETW infrastructure and Message Analyzer can capture their events. In this documentation, these are referred to as *system* ETW Providers, and in general, they write events from various applications and components on your system, such as Dynamic Host Configuration Protocol (DHCP), Domain Name System (DNS), Lightweight Directory Access Protocol (LDAP), and so on.

These system ETW Providers are accessible from the **Add System Provider** dialog, which you can display by clicking the **Add System Providers** item in the **Add Providers** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog during Live Trace Session configuration.

Filtering Events with Keywords and Level Settings
Most system ETW Providers from which Message Analyzer can capture event data are instrumented with filters that enable you to return specific provider events that you want to capture. The mechanism that allows you to do this is the provider event configuration which is specified in the provider manifest as **Keywords**, where each **Keyword** represents one of the provider's events. Message Analyzer enables you to select specific data from a Live Trace Session by specifying a low-level **Keyword** bitmask that matches the **Keyword** value of one or more system ETW Provider events, therefore effectively acting as a filter that tells the provider to write such events to the ETW session that enabled it. Another setting that you can configure for ETW Provider events is the error **Level**, which allows you to return events that correspond to a specified severity level, for example, **Critical**, **Error**, **Warning**, **Information**, and so on.

> **IMPORTANT**
> Some system ETW Providers may specify only a **Level** value for their events without also specifying a **Keyword** configuration for such events, and vice-versa. In addition, not all system ETW Providers that you add to your Live Trace Session configuration have event **Keyword** or error **Level** filtering available at all. In the latter case, this simply means that when the provider/manifest was created to instrument a particular Windows component with ETW technology, **Keyword** and/or **Level** filtering configurations were not specified by the developer. In this scenario, an ETW Provider will usually deliver all the triggered events from a particular component to the enabling ETW session, as there are no event **Keyword** and **Level** configurations specified for which a bitmask can filter.
>
> If event **Keyword** and error **Level** filter configurations are available for a chosen system ETW provider and you want to configure such filters, see the sections that follow to understand how to set and use them.

The conceptual section that follows provides some brief background information on event tracing to help clarify the meaning of these **Keyword** and **Level** filtering features.

Conceptual Background

Event tracing is built upon an API that exposes the following ETW components:

- **ETW Session** — provides an environment that accepts events, buffers them, and creates a trace file for logging the events or delivers them live in real-time to an ETW Consumer.
- **ETW Controller** — enables providers, starts and stops event tracing sessions, defines log files, obtains execution statistics, sets the buffer configuration, and so on. Note that a provider is turned on only when it is enabled for an ETW Session by the ETW Controller.
- **ETW Provider** — provides events to an event tracing session. A provider defines its interpretation of being enabled or disabled. In general, an enabled provider generates events, whereas a disabled provider does not.
- **ETW Consumer** — consumes the events from an event tracing session.

When an ETW Controller enables an ETW Provider, it exposes the provider event configuration to the ETW Session to enhance the provider's filtering instrumentation. An ETW Provider event configuration is specified with the use of the following two elements:

- **Level** — a 1-byte integer that enables filtering based on the severity or verbosity of events.
- **Keywords** — an 8-byte bitmask that enables the filtering of events from specific provider subcomponents.

For example, by selectively enabling these filtering features, the ETW Controller can enable providers to log the following:

- Only the error events from a particular provider subcomponent.
- All events from specific provider subcomponents.
- Specific events from provider subcomponents.

> **NOTE**
>
> When an ETW Controller enables a particular event **Level**, all provider events with a **Level** value that is less than or equal to what the ETW Controller specified are also enabled. For further details, see the Table ahead entitled "System ETW Provider Keyword and Level Filter Configurations".

Filtering with System ETW Provider Event Keywords and Levels

Message Analyzer provides the following filtering settings for system ETW Providers that appear in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog. The settings are accessible on the **ETW Core** tab of the **Advanced Settings** dialog that displays when you click the **Configure** link for any provider in the **ETW Providers** list:

- **Keywords(Any)** — enables you to specify an 8-byte, hexadecimal bitmask value that represents one or more **Keywords**, in order to set the category of events that an ETW event provider writes. The provider will write a particular event if the provider's event **Keyword** bits match *any* of the bits set in the **Keywords(Any)** bitmask.
- **Keywords(All)** — an optional 8-byte, hexadecimal bitmask that further restricts the category of events that an ETW event provider writes. If the provider **Keyword** for an event meets the **Keywords(Any)** condition, the provider writes the event only if the provider's event **Keyword** bits also match *all* of the bits set in the **Keywords(All)** bitmask. Note that this mask is not used if **Keywords(Any)** is set to zero.

> **NOTE**
>
> **Keyword** bitmask filters and provider event **Keyword** values are typically subjected to a logical ANDing process to validate whether an event should be written and returned to the ETW consumer (Message Analyzer) in an ETW session. For example, if the same bit positions in a provider event **Keyword** and the **Keywords(Any)** bitmask that you set are both logic 1, then the ANDing process signals that the associated event should be returned in the ETW session.

Using the ETW Keyword Filter Property Dialog

The figure that follows shows the **ETW Core** tab of the **Advanced Settings** dialog that you can display for any ETW Provider, as previously described. From this location, you can access **Keyword** bitmask configurations in the **ETW Keyword Filter Property** dialog, which displays after you click one of the ellipsis buttons (**...**) on the **ETW Core** tab of the **Advanced Settings** dialog. In the **ETW Keyword Filter Property** dialog, you can manually configure a **Keyword** bitmask, or you can select event **Keyword** presets to automatically create a **Keyword** bitmask.

In the figure, note that the **ETW Keyword Filter Property** dialog displays a **Keyword** bitmask value in the lower text box, which is the result of selecting several events in the **Automatic** list for the **Windows-NDIS-PacketCapture** provider, as indicated by check marks. Also note that you can view the decimal values of the listed events for this (or any) system ETW provider by using the *Windows Events Provider Explorer* tool described in Finding System ETW Provider Metadata. With this type of information, you can convert decimal values to hexadecimal values with an appropriate tool, should you ever need to manually configure **Keyword** bitmask settings.

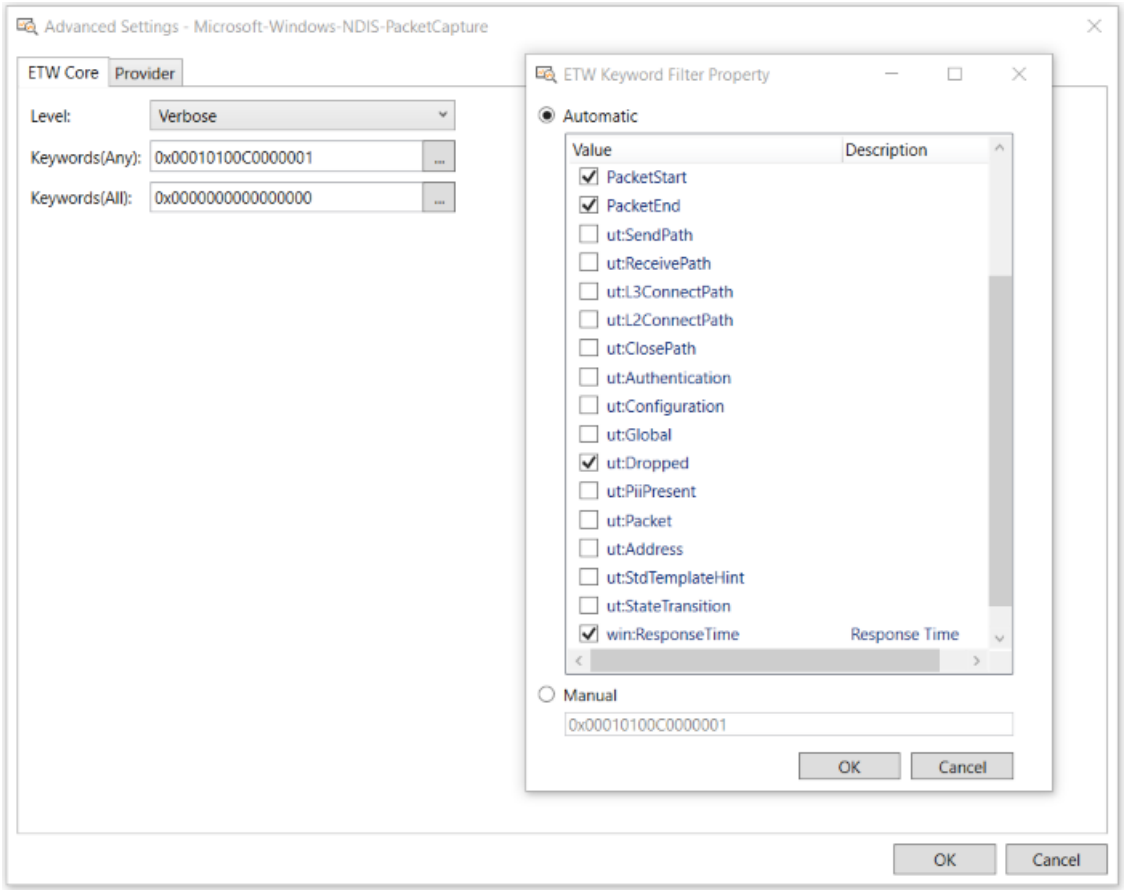

Figure 24: Advanced Settings Dialog ETW Keyword and Level Settings

Automatic vs Manual Keyword Bitmask Configuration

The easiest way to set a **Keyword** bitmask is to take advantage of the automatic configuration that is available in the **ETW Keyword Filter Property** dialog. However, you can also perform this process manually if you are familiar with (and can specify) the **Keyword** bitmask values that match the specific

events of an ETW provider. In either case, you should understand the meaning of the events that are distinguished by the provider **Keywords** for which you are specifying bitmasks, to enable a coherent interpretation of results. The subsections that follow provide some examples of manual **Keyword** bitmask configuration to return the events of a simple, but hypothetical provider. From this, you should be able to better understand what occurs during **Automatic** configuration, which is the feature you will likely use the most.

Manually Specifying a Keyword Bitmask Value

In many cases, if you leave the **Keywords(Any)** bitmask set to the default value of zero (0x0000000000000000) in the **ETW Keyword Filter Property** dialog, the associated ETW Provider will write all of its events, that is, those that are triggered. To return only specific events, you will need to set the **Keywords(Any)** bitmask to a value that will pass those provider events, which are typically defined in the ETW Provider manifest (see the System ETW Provider Event Configuration topic of the ETW Framework Conceptual Tutorial for details). To arrive at a value that will pass specific events, you will need to create a bitmask that is *inclusive* of all the provider events you want to return. The subsection that follows illustrates how a provider might define events with specific **Keyword** values and how you can return such events with a specific **Keyword** bitmask.

Provider Event Configuration — the table that follows shows some hypothetical **Keyword** definitions that could be specified in a manifest that is associated with an event provider.

Table 6. Example Provider Event Configuration

EVENT NAME	CONFIGURATION	BINARY VALUE	HEXADECIMAL VALUE
Initialization	Sets first binary bit.	000**1**	0x0000000000000001
File Read Operation	Sets second binary bit.	00**1**0	0x0000000000000002
File Write Operation	Sets third binary bit.	0**1**00	0x0000000000000004

Returning the Provider Events — in this hypothetical configuration of provider events, the inclusive binary value for all events written by this provider is 0x0000000000000007 (equivalent to 0**111** binary, or 7 in decimal). If you wanted to receive all of the above provider events, you would set the **Keywords(Any)** bitmask in the Message Analyzer **ETW Keyword Filter Property** dialog to a hex value of 0x0000000000000007. However, if you were interested in receiving Initialization and File Read Operation events only, you would set the **Keywords(Any)** bitmask in the **ETW Keyword Filter Property** dialog to a hexadecimal value of 0x0000000000000003 (equivalent to 00**11** in binary, or 3 in decimal).

Note that in the earlier scenario with the **Keywords(Any)** bitmask set to 0x0000000000000007, any of the above specified provider events will be returned, if they are triggered. However, if you also set the **Keywords(All)** bitmask to a hex value of 0x0000000000000004 (**0100** binary), this restriction would allow only the File Write Operation events to be returned, if they are triggered, as described earlier in this section.

Setting the Error Level

For events that are delivered to an enabling ETW session, you can obtain an indication of the severity or verbosity of event errors by setting the error **Level** that the session will report for the events of a particular ETW Provider. To do this, select a particular value from the **Level** drop-down list on the **ETW Core** tab of the **Advanced Settings** dialog for a any ETW provider. The values that you can set are described in the table that follows, along with some further details about **Keyword** configuration options.

TIP

You may be able to discover event **Keyword** and error **Level** settings for various trace providers on your system by referring to the topic Finding System ETW Provider Metadata.

Table 7. System ETW Provider Keyword and Level Filter Configurations

CONFIGURATION SETTING	VALUES	DESCRIPTION
Level	You can configure this setting to one of the following values: - **LogAlways** (0x0) - **Critical** (0x1) - **Error** (0x2) - **Warning** (0x3) - **Information** (0x4) - **Verbose** (0x5)	Specifies the level of detail included in the ETW provider event. Levels indicated in the **Values** column to the left are inclusive. For example, if you set the **Level** to **Verbose**, the provider will write all **Critical**, **Error**, **Warning**, and **Information** events as well. If you set the **Level** to **Warning**, the provider will also write all **Critical** and **Error** events. Note that if you set the **Level** to **LogAlways**, it ensures that all error events will always be written.
Keywords(Any)	You can configure this setting in either of the following ways: - **Manual** — you can manually specify an 8-byte hexadecimal bitmask value to enable a system ETW Provider to write events whose **Keyword** bits match *any* of the bits in the specified **Keyword (Any)** bitmask. - **Automatic** — you can select one or more preset keywords to automatically configure an 8-byte hexadecimal bitmask value, to enable a system ETW Provider to write events whose **Keyword** bits match *any* of the bits in the specified **Keyword (Any)** bitmask.	Provides a convenient way to add filtering at the kernel level, which enhances performance as follows: - The provider selects specific data to retrieve, thereby reducing the number of messages being captured, which subsequently increases the speed at which data is captured. - Filtering at kernel level is inherently faster than user mode filtering (following the parsing process). You can set a **Keywords(Any)** bitmask filter by clicking the ellipsis [**...**] to the right of the current hexadecimal **Keywords(Any)** value, to display the **ETW Keyword Filter Property** dialog. From this dialog, you can either select **Manual** to specify a **Keywords(Any)** bitmask value, or you can select **Automatic** to choose a bitmask based on a preset bitmask **Value**, which indicates a subcomponent of the provider. **Note:** Before setting this value, you should be familiar with the **Keyword** settings of the provider event for which you are trying to filter. You may obtain some of this information by consulting the system ETW Provider manifest, as described in Finding System ETW Provider Metadata.

CONFIGURATION SETTING	VALUES	DESCRIPTION
Keywords(All)	You can configure this setting in either of the following ways: - **Manual** — you can manually specify an 8-byte, hexadecimal bitmask value, to enable a system ETW Provider to write events whose **Keyword** bits match *all* of the bits in the specified **Keyword (All)** bitmask. - **Automatic** — you can select one or more preset keywords to automatically configure an 8-byte, hexadecimal bitmask value, to enable a system ETW Provider to write events whose **Keyword** bits match *all* of the bits in the specified **Keyword (All)** bitmask.	Provides a convenient way to add filtering at the kernel level, which enhances performance as described above. You can set a **Keywords(All)** bitmask by clicking the ellipsis [**...**] to the right of the current hexadecimal **Keywords(All)** value, to display the **ETW Keyword Filter Property** dialog. From this dialog, you can either select **Manual** to specify a **Keywords(All)** bitmask value, or you can select **Automatic** to choose a bitmask based on a preset **Value**, which indicates a subcomponent of the provider. As indicated earler, using this filter further restricts the events that will be written by the system ETW Provider. Only if a provider's event **Keyword** matches the **Keywords(Any)** condition and only if all bits in the **Keywords(All)** bitmask also exist in the provider's event **Keyword** configuration will the provider write the specific event/s.

Finding System ETW Provider Metadata

This topic describes several possible ways to locate the event **Keywords** that a system ETW Provider writes. These methods are discussed in the subsections that follow.

Opening the ETW Keyword Filter Property Dialog

You can view predefined event **Keyword** configurations in the **ETW Keyword Filter Property** dialog that displays when you click the ellipsis on the **Keywords(Any)** or **Keywords(All)** drop-downs that appear on the **ETW Core** tab of the **Advanced Settings** dialog for any particular ETW provider. Although, you may need to do some research to discover the meaning of the events.

Using the WEPExplorer Graphic Utility

Each system ETW Provider that is installed and registered on your system contains metadata that is stored in a manifest. To locate this metadata, which includes event **Keywords**, error **Levels**, **Opcodes**, **Channels**, and so on, you might consider using a graphic utility such as the *Windows Events Provider Explorer* (WEPExplorer) to obtain this information, as described in Windows Events Provider Explorer on the web. If you are a programmer, you can also return this information programmatically, as described in Getting a Provider's Metadata on MSDN.

Opening Performance Monitor

You also might be able to review the **Keyword** configuration of a system ETW Provider for which you are looking by opening the **Performance Monitor** tool and locating the provider as it running in a live Windows Event Tracing Session. To view the event **Keyword** and error **Level** configuration for events that such a provider writes, follow the steps below.

1. From the **Start** menu or from the desktop, right-click **Computer** or the **Computer** icon, respectively, and select the **Manage** item to display the **Computer Management** console.

2. In the **Computer Management (Local)** pane, expand the **Performance** node, expand the **Data Collector Sets** node, and then click **Event Trace Sessions**.

   The name and status of event trace sessions that are running on your machine are displayed.

3. Right-click an event trace session such as **EventLog-System** and select the **Properties** item from the menu.

   The **EventLog-System Properties** dialog displays. This may take a few moments.

4. Select the **Trace Providers** tab and then double-click a provider in the **Providers** list box.

   The current **Keyword** and **Level** configuration of the provider you clicked displays in the **Properties** list box.

5. In the **EventLog-System Properties** dialog, click the **Edit** button to display a list of all the **Keywords** that define the events that the selected provider can write to a trace consumer, which are typically specified in the provider manifest.

See Also

Generating a Provider Manifest

Setting the Session Focus

2 minutes to read

This section describes various high-level settings that you can specify for a Live Trace Session that alter the overall focus of the session. Although low-level provider filter settings can create the focus on capturing specific messages, you have the option to specify several high-level settings that can further refine the focus of data capture, improve the efficiency of data capture mechanisms, and specify how you will view results. For example, by selecting a high-level **Session Filter** or a **Parsing Level**, you can refine the session focus by specifying the type of data you want to capture, or inversely, the traffic you want Message Analyzer to ignore by filtering it out.

However, before you specify either of these filtering elements, you might want to ensure that you do not have any low-level provider filter settings that conflict with the high-level filter settings that you specify. Another thing to consider is that low-level provider filters such as **Fast Filters**, **WFP Layer Set** filters, and event **Keyword** bitmask filters are all typically highly performant filters that consume less system resources, while a **Session Filter** requires additional parsing time. On the other hand, by setting a **Parsing Level** that limits the upper layer to which Message Analyzer will parse messages, you can also realize performance improvements. Moreover, when specifying a **Session Filter**, you can take advantage of numerous built-in Filter Expressions that provide a wide variety of functionality and you can easily change or remove the effects of a **Session Filter** to facilitate different analysis perspectives.

You can also specify several advanced settings to optimize the underlying ETW Session in which Message Analyzer captures events. Other session-level configuration that you can perform is the setup for decrypting TLS- and SSL-encrypted messages. Lastly, you have the option to set the focus in which you will view the session results data, by choosing one of the default data viewers prior to starting a Live Trace Session.

These discussions are covered in the following subtopics:

Selecting Data to Capture
[Specifying Advanced ETW Session Configuration Settings]pecifying-advanced-etw-session-configuration-settings.md)
Decrypting TLS and SSL Encrypted Data
Selecting a Session Data Viewer

See Also

Modifying Default Provider Settings

Selecting Data to Capture

8 minutes to read

As part of the browse-select-view (BSV) model, Message Analyzer provides a *data selection* feature that enables you to define the scope of the information that you capture or load through a session. Although the concept of data selection applies equally to capturing data in a Live Trace Session or loading saved data through a Data Retrieval Session, with exception of using different filter types in some cases, this section focuses on selecting data in a Live Trace Session.

Using Data Selection

Data selection is not exactly a specific user interface element that you can locate, but rather, more of an approach you can take to focus on obtaining the precise data you need to work with to quickly solve problems. In Message Analyzer, the goal of data selection is to acquire the least amount of data necessary to resolve an issue, in order to minimize consumption of system resources, improve performance, and streamline the data analysis process.

Message Analyzer provides several tools that you can use to create this focus, as indicated in the list below. To use these data selection capabilities, you will typically create a Live Trace Session (or Data Retrieval Session) that selects specific data based on configurable criteria, prior to starting the actual data capture process. During data capture, the selection criteria that you configured is applied. This context enables you to narrow the focus of the data capture process to only the message data that you want to work with.

Isolating Specific Types of Message Data

A very effective means of isolating specific message types in a Live Trace Session consists of using filters to return only the type of message data that you choose, while blocking all data that does not specifically meet your designated filtering criteria. The effects of filtering can be initiated at a high-level or a low-level to select the data that you want to extract from a Live Trace Session, as described in the subsections that follow.

Selecting Data with High-Level Filtering
High-level filtering enables you to alter the overall session results that you obtain and will also impact the effects of any low-level filters that you specified. As a result of applying such a combination of filter types, you might not return the results you expect unless you consider how to combine them correctly. For example, if you specify a low-level **Fast Filter** that passes specific traffic on a particular port, you should ensure that you do not also specify a high-level **Session Filter** that somehow overrides the effects of the specified **Fast Filter**, or you might fail to see the data you are expecting. However, note that while you have the option of changing or removing a **Session Filter** in the **Edit Session** dialog for a set of trace results, you cannot alter the effects of an applied **Fast Filter**.

You might also consider that you can limit the level to which Message Analyzer parses by specifying a **Parsing Level**. For example, if you specified a low-level **Fast Filter** to pass some Application Layer traffic on a particular port and you also limit parsing to the Network Layer, Message Analyzer will not parse or display any of the Application Layer traffic.

During configuration of a Live Trace Session, you can specify the primary high-level filters as follows. Note that a **Session Filter** and the **Parsing Level** drop-down list are shown in the figure of the topic Using a Session Filter.

- **Session Filter** — select a **Session Filter** item from the **Message Analyzer Filters** asset collection **Library** in the **Session Filter** text box of the **New Session** dialog, in a category that is relevant to the **Trace Scenario** you have selected for your Live Trace Session. For example, if you selected the **Pre-Encrypted for HTTPS** scenario from the **Select Scenario** drop-down list on the **Live Trace** tab in the **New Session** dialog, it could be advantageous to also select the `HTTP.StatusCode >=400` **Session Filter** in

the **HTTP** category to focus on HTTP errors that are occurring on the client.

For further details about how to select data from a Live Trace Session with the use of a **Session Filter**, see Working with Session Filters in a Live Trace Session.

- **Parsing Level** filter — choose a **Parsing Level** to return a set of messages that are constrained by the upper stack level to which Message Analyzer parses. This feature simultaneously creates a unique analysis perspective and a focused set of messages, while improving performance by removing all messages above the specified **Parsing Level**.

  For further details about how to select data from a Live Trace Session with the use of a **Parsing Level**, see Setting the Session Parsing Level.

> **TIP**
> You can also filter trace *results* data by applying a view **Filter** in an Analysis Session. Note that while you can remove the effects of a view **Filter** on a set of trace results, the effect of using a **Fast Filter**, as described ahead, is inherent to a given set of trace results and cannot be changed unless you rerun the trace with a change in the **Fast Filter** configuration. However, you have the option to modify a **Session Filter** or specify a different one and apply those changes to a set of trace results, if you do so through the **Edit Session** dialog, which is accessible only after your initial trace completes.

Selecting Data with Low Level Filtering

Message Analyzer also enables you to specify other types of filtering during Live Trace Session configuration, such as low-level **Fast Filters**, **WFP Layer Set** filters, event **Keyword** bitmask and error **Level** filters, and others, You can even use **Trace Scenarios** to select specific data in a Live Trace Session. These filters consist of the following:

- **Fast Filter** — configure these filters from the **Advanced Settings** dialog for **Trace Scenarios** that use the **Microsoft-PEF-WFP-MessageProvider** or the **Microsoft-PEF-NDIS-PacketCapture** provider. **Fast Filters** work at the provider/driver level, which means that you can isolate the messages you want to focus on without incurring any Message Analyzer parsing time to apply the filtering criteria. This is partly what makes **Fast Filters** so efficient and quick.

  For further details about how to select data from a Live Trace Session with the use of a **Fast Filter**, see the PEF-NDIS Fast Filters and PEF-WFP Fast Filters topics.

- **WFP Layer Set** filters — use these filters to specify the direction in which data is captured at the Transport Layer for IPv4 and IPv6, in scenarios that use the **Microsoft-PEF-WFP-MessageProvider**.

  For further details about how to select data from a Live Trace Session with the use of **WFP Layer Set** filters, see the PEF-WFP Layer Set Filters topic.

- **WebProxy** filters — use HTTP filters such as **Hostname** and **Port** to isolate HTTP messages from a particular host or on a specified port in scenarios that use the **Microsoft-PEF-WebProxy**/Fiddler provider, such as the **Pre-Encrypted for HTTPS** scenario.

  For further details about how to select data from a Live Trace Session with the use of a **WebProxy** filter, see the WebProxy Filters topic.

- **Keyword** and **Level** filters — specify event **Keyword** bitmasks and/or error **Level** settings to match system ETW Provider event configurations and thereby select specific events to be returned from a Live Trace Session.

  For further details about specifying **Keyword** bitmasks and **Level** settings, see System ETW Provider Event Keyword/Level Settings.

- **Adapter** filters — select specific data from a trace by applying **Fast Filter** groups to local host adapters in

scenarios that use the **Microsoft-PEF-NDIS-PacketCapture** provider; or apply remote host adapter or Hyper-V-Switch layer, payload, and other special filters in scenarios that use the **Microsoft-Windows-NDIS-PacketCapture** provider.

For further details about how to select data from a Live Trace Session with the use of an Adapter filter during **Microsoft-PEF-NDIS-PacketCapture** provider configuration, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.
For further details about how to select data from a Live Trace Session with the use of an Adapter filter during **Microsoft-Windows-NDIS-PacketCapture** provider configuration, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

Selecting Data with a Trace Scenario
The **Trace Scenario** that you select for your Live Trace Session can have an impact on the data that you return. For example, you could select a **Trace Scenario** that uses a message provider that focuses on a particular stack level, such as the **Loopback and Unencrypted IPSEC** scenario. This scenario focuses on messages above the IP/Network Layer while filtering out most lower-level noise at the Data Link Layer, such as broadcast traffic.

You could also select a **Trace Scenario** that employs predefined filtering to return a specific result, such as the **Local Loopback Network** scenario, which uses two **Fast Filters** and **WFP Layer Set** filters to pass only loopback traffic. Moreover, you might select a **Trace Scenario**, such as the **VPN** scenario, that employs a number of system ETW Providers that focus on returning specific messages and events that are useful when you are troubleshooting VPN issues.

You might also create your own **Trace Scenario** by selecting one or more system ETW providers from the ***Add Provider*** drop-down list in the **New Session** dialog during Live Trace Session configuration, to create focus on specific types of messages to retrieve during live capture.

For further details about selecting a **Trace Scenario** to return specific types of data, see the Selecting a Trace Scenario section.

More Information
To learn more about the configuration features of the **Advanced Settings** dialogs for the **Microsoft-PEF-WFP-MessageProvider**, **Microsoft-PEF-NDIS-PacketCapture** provider, and the **Microsoft-Windows-NDIS-PacketCapture** provider, see the following topics:
Using the Advanced Settings- Microsoft-PEF-WFP-MessageProvider Dialog — describes how to configure **Fast Filters**, **Discarded Packet Events**, and **WFP Layer Set** filters (packet direction at the Transport Layer).
Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog — describes how to select adapters, configure **Fast Filters**, and create **Fast Filter** groups to apply to local host adapters.
Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog — describes how to configure advanced packet direction and layer filtering configurations for the NDIS stack and Hyper-V-Switch extension layers on remote (or local) hosts, in addition to packet **Truncation** filters and special message payload filtering such as **EtherType** and **IP Protocol Numbers**.

See Also

Editing Existing Sessions

Working with Session Filters in a Live Trace Session

5 minutes to read

Prior to starting any Live Trace Session, you can specify a **Session Filter** that will be applied by the Runtime in *user mode* to messages captured by one or more message providers that are included in the **Trace Scenario** you selected. To specify a **Session Filter**, you can select one of the built-in filters from the centralized **Message Analyzer Filter** asset collection **Library**. This user **Library** is accessible from the following locations during the indicated session phases, although you can only use the first location below to specify a **Session Filter**:

> **NOTE**
>
> You can also create your own Filter Expression in either of these same locations and save it to the **Library**, at which point you can use it as a **Session Filter**.

- **Session configuration** — locate a **Session Filter** in the centralized **Message Analyzer Filter** asset collection **Library** drop-down list above the **Session Filter** text box in the **New Session** or **Edit Session** dialog.
- **Session analysis** — locate a view **Filter** from the same **Library** on a Filter panel of the Filtering toolbar that appears above the analysis surface of any data viewer in a set of trace results.

> **NOTE**
>
> You can also apply any Filter Expression item in the **Message Analyzer Filter** asset collection **Library** as a **Session Filter** when configuring a Data Retrieval Session. This **Library** is in the identical location as when you are configuring a Live Trace Session in the **New Session** dialog.

> **IMPORTANT**
>
> Although a **Session Filter** can reduce the noise of unwanted traffic and narrow the scope of data capture, Message Analyzer must still parse messages according to the **Session Filter** criteria, which can take time. As a result, there is the possibility that messages could be dropped in very heavy traffic conditions. While you may be able to offset this possibility by configuring **Live Trace Message Buffer** settings (click the **Options** item in the Message Analyzer **Tools** menu and go to the **General** tab of the **Options** dialog), there are limitations on how much you can compensate.

Selecting Built-In Session Filters

To select a built-in Filter Expression for a Live Trace Session, click the **Library** drop-down list on the toolbar above the **Session Filter** text box in the **New Session** dialog and then select a filter item. After you select a filter, the filter code displays in the **Session Filter** text box. Note that you can optionally modify any built-in **Session Filter** configuration *for the session only*, by specifying your Filter Expression changes prior to starting your Live Trace Session. Note that you cannot alter the default configuration of any built-in Filter Expression item that is contained in the centralized Filter Expression **Library**. The built-in **Library** Filter Expressions that are included by default in Message Analyzer fall into the following categories of application. The functions of the filters that exist in the below categories are described in Filtering Live Trace Session Results, with exception of the **Azure Storage** category:

> **NOTE**
>
> The **Azure Storage** category filters are not described in this Operating Guide, as Microsoft provides related information in Azure blogs and tutorials, which you can access from the topic Filtering Live Trace Session Results. Note that the **Azure Storage** category will only exist in your user **Library** if you download the **Azure Storage Filter** asset package with the use of the **Asset Manager** dialog.

- **Azure Storage**
- **Address Filtering**
- **Diagnosis**
- **General Examples**
- **RegEx**
- **Contains Filters**
- **HTTP**
- **TCP**
- **LDAP**
- **Remove Noise**
- **File Sharing**
- **USB**

Creating New Session Filters

To create your own **Session Filter**, you can write a Filter Expression in the **Session Filter** text box of the **New Session** or **Edit Session** dialogs prior to starting a Live Trace Session or prior to applying configuration changes to an existing session, respectively. You can then click the **New Filter** item in the **Library** drop-down list to display the **Edit Filter** dialog, which captures the filter code that you specified and enables you to **Save** the new Filter Expression to the **Examples** category of your user **Library**.

However, before you write a new Filter Expression, you might want to review the topics specified in **More Information** to familiarize yourself with some of the built-in filters provided with Message Analyzer and the Filtering Language with which they were created. Also, to assist you in Filter Expression development, Message Analyzer provides the Filter IntelliSense Service, which activates as soon as you begin typing in the **Session Filter** text box. You can also open the **Field Chooser Tool Window** and navigate through the message hierarchy of supported protocols to see a tree list view of such hierarchies, so that you can discover at a glance the defined fields that are accessible for your Filter Expressions.

You can also create and save a view **Filter** in an Analysis Session, which is then added to the centralized **Message Analyzer Filter** asset collection **Library**; at which point you can select that filter from the **Library** on the toolbar above the **Session Filter** text box of the **New Session** or **Edit Session** dialog to add it as a **Session Filter** to a current or subsequent Live Trace Session or Data Retrieval Session.

To create a new Filter Expression from any of the previously indicated locations, you must select the **New Filter** item in the centralized **Library** drop-down list. Thereafter, the **Edit Filter** dialog displays, from where you can create a new Filter Expression, either by developing a new filter configuration or by modifying an existing/predefined Filter Expression from the **Library**.

NOTE

If you have created a Filter Expression that uses an **Alias** (typically a friendly name that replaces some cryptic field value), that filter will appear in your user **Library** drop-down list in the **New Session** dialog. You can use such a Filter Expression that contains an **Alias** just as you would any other filter.

Applying a Session Filter

As described in several topics in this Operating Guide, adding a **Session Filter** to your Live Trace Session configuration enables you to focus on capturing specific data of interest after you have started a Live Trace Session. **Session Filters** also enable you to reduce message count and realize better performance. When you start a Live Trace Session, the messages that are captured by one or more providers in the **Trace Scenario** that you selected for the session are all automatically subjected to the filtering criteria of the **Session Filter** that you specified.

NOTE

When you apply a **Session Filter** to a Live Trace Session, a funnel icon displays to the right of the corresponding top-level session node in the **Session Explorer Tool Window**, to indicate that the session has had a **Session Filter** applied to it. This icon provides a quick reminder of the initial configuration status of the session.

More Information

To learn more about using the Filtering Language, see Writing Filter Expressions.

To learn more about the built-in filters in the centralized **Message Analyzer Filter** asset collection **Library** that are provided with Message Analyzer, see Filtering Live Trace Session Results.

To learn more about using the statement completion feature for Filter Expressions, see the Filter IntelliSense Service topic.

To learn more about the **Field Chooser**, see the Field Chooser Tool Window topic.

See Also

Filtering Message Data

Managing Session Filters

6 minutes to read

Session Filter management features enable you to create, modify, delete, import, export, and share filters. These same management features apply to all Message Analyzer asset collections that you typically work with, such as **Color Rules**, view **Filters**, viewer **Layouts**, **Viewpoints**, **Trace Scenarios**, **Chart View Layouts**, **Pattern Expressions**, and so on.

Using the Edit Filter Dialog

When you are configuring a Live Trace Session, you have the option to select or create a **Session Filter**, as described in Working with Session Filters in a Live Trace Session. To create and save a new Filter Expression that you can use as a **Session Filter** or view **Filter**, you must open the **Edit Filter** dialog by selecting the **New Filter** item from the **Library** drop-down list above the **Session Filter** text box in the **New Session** dialog. In a similar manner, you can also access the **Edit Filter** dialog from the **Library** in any Filter panel that displays on the Message Analyzer Filtering toolbar during an Analysis Session to create a new Filter Expression. However, you can also create a **Session Filter** from the **Manage Filter** dialog by modifying a copy of an existing Filter Expression; the **Manage Filters** dialog is accessible by clicking the **Manage Filters** item in the **Library** drop-down list in either of the indicated locations. In either case, you end up displaying the same **Edit Filter** dialog; the only difference is in how you get there. The **Edit Filter** dialog is shown in the figure that follows:

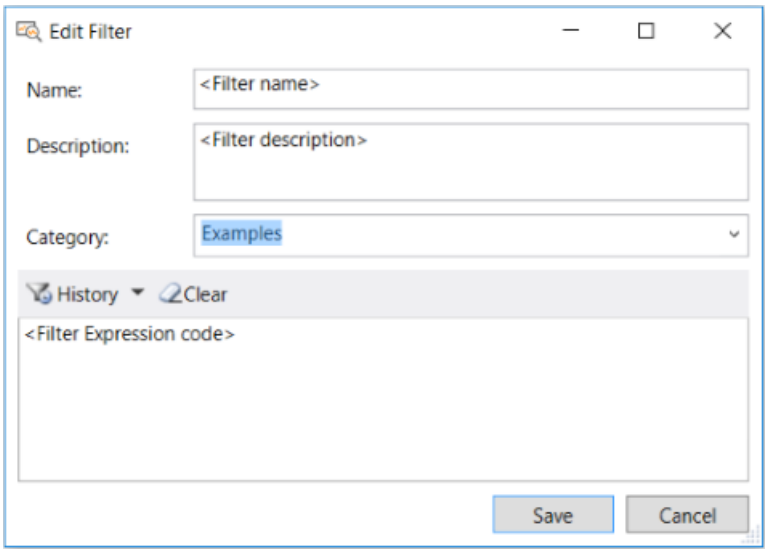

Figure 25: Edit Filter dialog

For example, when you select the **New Filter** item in the view **Filter** or **Session Filter** locations of the central Filter Expression **Library**, the **Edit Filter** dialog immediately displays and enables you to create and save a new Filter Expression. But when you select the **Manage Filters** item in either **Library** drop-down list, the **Manage Filter** dialog displays first; you can then create a new Filter Expression by right-clicking any of the built-in filter items in the **Manage Filter** dialog and selecting the **Create a Copy** command in the context menu that appears. Thereafter, the **Edit Filter** dialog displays from where you can modify the filter configuration, rename it, and save it back to the centralized Filter Expression **Library** as a new filter. Note that the **Edit** command is available for any Filter Expression that you right-click in the **My Items** category of the **Manage Filter** dialog, that is, if you want to modify one of your own custom-designed filters. Otherwise, you cannot edit any of the built-in Filter Expressions.

After configuring a new **Session Filter**, you can then select it during Live Trace Session or Data Retrieval Session configuration. You will also be able to select your new filter from the **Library** drop-down list on the Message Analyzer Filtering toolbar during an Analysis Session, but in this case, it will function as a view **Filter** only.

NOTE

Creating Filter Expressions from the **Manage Filter** dialog for a Live Trace Session is only one use of this feature, as the dialog has additional management features that are briefly described in the remainder of this section.

Centrally Managing Filters

The **Library** in which **Session Filters** are contained is a local, centralized Filter Expression **Library** that is populated by the **Message Analyzer Filters** asset collection, which is accessible from the **Asset Manager** dialog. You can manage the Filter Expressions from this collection in your user **Library** by using the **Manage Filter** dialog. Because Filter Expressions are contained in a single central **Library**, you can not only centrally manage the **Library**, but you can also use any Filter Expression that is contained in the central **Library** as a **Session Filter** when configuring a Live Trace Session or a Data Retrieval Session, or as a view **Filter** during a data Analysis Session. You can manage the **Library** from either the toolbar above the **Session Filter** text box of the **New Session** (or **Edit Session**) dialog, or from the Message Analyzer Filtering toolbar that appears in every Analysis Session. In both cases, you access the **Manage Filter** dialog by selecting the **Manage Filters** item from the **Library** drop-down list in these locations. The **Manage Filter** dialog is shown in the figure that follows:

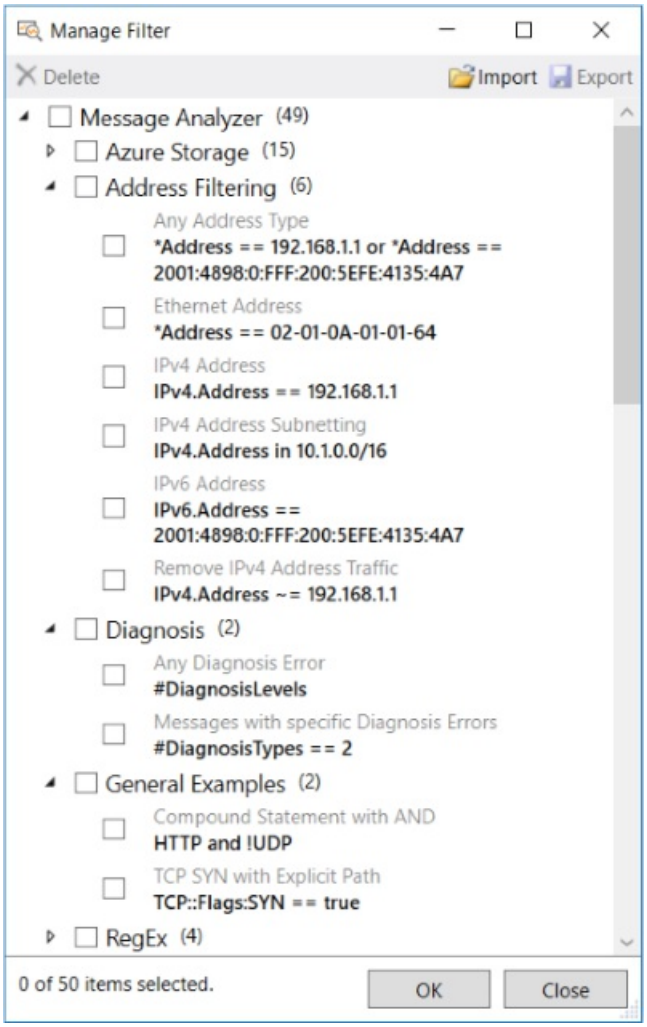

Figure 26: Manage Filter dialog

Message Analyzer uses the same management dialog format to manage all user Library types; the only difference is the varying types of assets that you manage, for example **Color Rules**, **Filters**, **Trace Scenarios**, **Viewpoints**, viewer **Layouts**, **Pattern Expressions**, **Chart Viewer Layouts**, and so on. Therefore, in this Operating Guide, the management dialog is occasionally referred to generically as the **Manage <AssetType>** dialog. From the **Manage <AssetType>** dialog, you can import specific items; or export, delete, and modify selected items such as

Filter Expressions; as described in Managing User Libraries.

Sharing Filter Items

The centralized Filter Expression **Library** items are shareable, as are items from other Message Analyzer asset collections. Message Analyzer provides a simple way to expose the Filter Expressions contained in this **Library** to others for sharing, or to retrieve Filter Expressions that others have shared.

To share Filter Expression items directly with others, you can use the **Export** feature in the **Manage Filter** dialog to save one or more Filter Expression items to a designated file share or other location. In addition, you can use the **Import** feature in the same dialog to access Filter Expression items that have been shared by others at a designated location. When exporting Filter Expressions, you have the option to select specific items that you want to distribute to others, including any items that you have created or modified. When importing a Filter Expression asset collection, you can choose the items you want to retrieve.

You can also share your Filter Expression items through a user-configured feed in the Message Analyzer Sharing Infrastructure. You can create such a feed from the **Settings** tab of the Message Analyzer **Asset Manager** (accessible from the global Message Analyzer **Tools** menu), by clicking the **Add New Feed** button and specifying a feed name and directory location. Thereafter, you can use the **Export** feature of the **Manage Filter** dialog to post your Filter Expression items as an asset collection to the configured feed location. You can also update an existing asset collection and make it available to team members or other users through your configured feed, where they can view, synchronize, and download the collection. However, for users to synchronize with collection updates, some manual configuration is necessary in the current Message Analyzer release, as described in Manual Item Update Synchronization. In future Message Analyzer releases, the Sharing Infrastructure publishing features may automatically enable others to synchronize with asset collection updates when they are shared through user feeds.

Accessing the Message Analyzer Filters Asset Collection

Microsoft provides a default **Message Analyzer** feed on the **Downloads** tab of the **Asset Manager** that enables you to download the **Message Analyzer Filters** asset collection once from a Microsoft web service and to synchronize with collection updates that are periodically pushed out by the service. To receive updates that will appear in the **Message Analyzer** category of the centralized Filter Expression **Library**, you must set the **Message Analyzer Filters** asset collection to the auto-sync state from the **Downloads** tab of the Message Analyzer **Asset Manager**, that is, if the collection is not already synced. At any time after setting the auto-sync state, you can perform a download of an auto-synced collection from the **Settings** tab of the **Asset Manager**. The **Message Analyzer Filters** asset collection is shown on the **Downloads** tab of the **Asset Manager** dialog in the figure that follows:

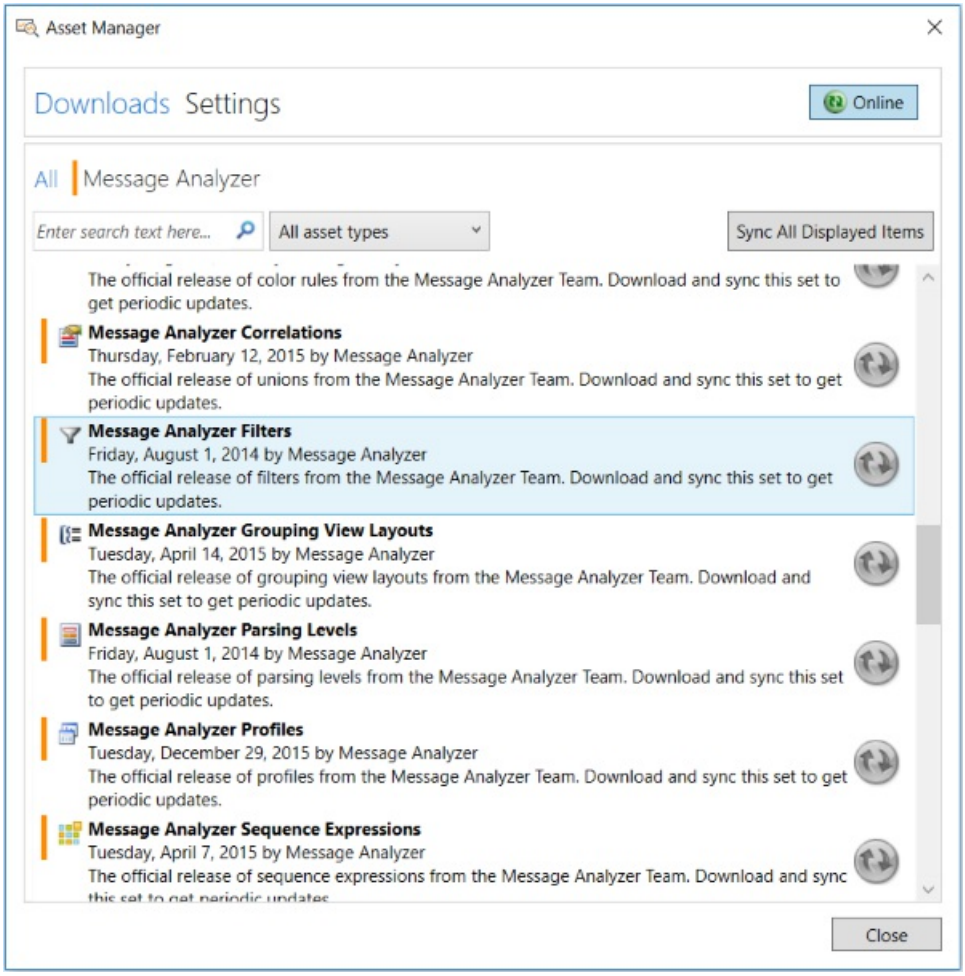

Figure 27: Asset Manager dialog

> **NOTE**
>
> If an automatic update to an auto-synced asset collection is pushed out by the web service to your Message Analyzer installation, the updates are automatically populated to your user **Library**, with no further action required.

More Information

To learn more about the Message Analyzer Sharing Infrastructure, including how to share asset collections with others and to auto-sync asset collections for updates, see the Managing Message Analyzer Assets section.

To learn more about the common **Manage <AssetType>** dialog, see Managing User Libraries.

Setting the Session Parsing Level

9 minutes to read

Message Analyzer provides a collection of built-in **Parsing Level** scenarios that are available for selection in the **New Session** or **Edit Session** dialog. You can specify a single **Parsing Level** for any one particular session only. By carefully choosing a **Parsing Level** that is appropriate for the **Trace Scenario** you are using, as described in Choosing Parsing Levels, you can dramatically improve your Message Analyzer experience because you will be able to do the following:

- **Realize performance improvements when loading large traces** — by limiting the level to which Message Analyzer parses, you can quicken processing time, consume less memory, and reduce the overall costs of problem solving.
- **Target specific messages for analytical focus** — by removing messages above a specified stack layer or module, you can more easily focus on messages at a layer of interest.

Understanding Parsing Levels

To facilitate the previously mentioned improvements, each **Parsing Level** scenario applies preconfigured Runtime filtering functionality that defines the layer or module to which Message Analyzer parses, and specifies other messages that are explicitly passed in each scenario. For example, a typical predefined filter has a StopAtModule function that tells the Runtime to stop at a particular module such as TCP; it can also have an ExcludeFilter that enables certain other message types to pass that were predetermined to make specific **Parsing Level** scenarios more useful for diagnostic analysis at the particular layer where they apply.

The following table describes the **Parsing Levels** that you can select from the **Parsing Level** drop-down menu in the **New Session** or **Edit Session** dialog. Note that the **Definition** column specifies a representation of the filtering functionality that is preconfigured in each **Parsing Level** scenario.

Table 8. Message Analyzer Parsing Levels

NAME	DESCRIPTION	DEFINITION
Full	Enables the Message Analyzer default **Parsing Level**.	No **Parsing Level** is applied; note however, that parsing results can be impacted by the PEF **Trace Scenario** or other message provider in use.
Network Analysis	Focuses on diagnosing Transport and Network layers, but includes UDP/TCP traffic along with DNS, DHCP, ARP, and ICMP messages.	- Exclude Filter: TCP.Port==53 or TCP.Port==42 StopAtModule: TCP - Exclude Filter: UDP.Port==53 or UDP.Port==546 or UDP.Port==67 or UDP.Port==137 or UDP.Port==1512 StopAtModule: UDP **Note:** ARP and ICMP are included in this scenario because they are not explicitly excluded, therefore, they parse to the top level in this scenario.

NAME	DESCRIPTION	DEFINITION
File Sharing	Focuses on diagnosing file sharing scenarios, but is limited to SMB and some simple name resolution protocols such as WINS, DNS, and NetBIOS.	- Exclude Filter: TCP.Port==445 or TCP.Port==135 StopAtModule: TCP - Exclude Filter: StopAtModule: SMB - Exclude Filter: StopAtModule: SMB2 - Exclude Filter: UDP.Port == 53 or UDP.Port==546 or UDP.Port==67 or UDP.Port==137 or UDP.Port==1512 StopAtModule: UDP
High Performance Capture without Parsing	Displays events at the ETW level only, for the highest possible performance.	- Exclude Filter: StopAtModule: ETW - Exclude Filter: StopAtModule: CapFile - Exclude Filter: StopAtModule: PcapFile
HTTP	Focuses on HTTP traffic.	- Exclude Filter: IPv4.Protocol==6 StopAtModule: IPv4 - Exclude Filter: IPv6.NextHeader==6 StopAtModule: IPv6 - Exclude Filter: WFPCapture.Protocol==6 StopAtModule: WFPCapture - Exclude Filter: TCP.Port==80 StopAtModule: TCP - Exclude Filter: StopAtModule: HTTP
Identity and Active Directory	Focuses on LDAP traffic.	- Exclude Filter: IPv4.Protocol==6 StopAtModule: IPv4 - Exclude Filter: IPv6.NextHeader==6 StopAtModule: IPv6 - Exclude Filter: TCP.Port==389 or TCP.Port==3268 StopAtModule: TCP

> **NOTE**
> Because upper message layers are removed by the application of these **Parsing Level** scenarios, which subsequently reduces the number of messages that Message Analyzer parses, you should notice incremental but very striking performance gains.

Choosing Parsing Levels

Because Message Analyzer PEF providers focus on different inspection points into the message stack, you should carefully consider the **Parsing Level** that you choose; otherwise, you might obtain unexpected results. For example, you might choose a **Loopback and Unencrypted IPSEC Trace Scenario** for your Live Trace Session to reduce lower-level noise and to focus on Transport layer messages, but you are also interested in looking at messages from a particular application along with their transports. If you also choose the **Network Analysis Parsing Level**, Message Analyzer will parse up to and including the Network Layer only, which means that you will not see the Application layer traffic that the **Loopback and Unencrypted IPSEC** scenario would normally pass.

On the other hand, if you were to choose the **Pre-Encryption for HTTPS Trace Scenario** and **Network Analysis Parsing Level** scenario to work together in a trace, you would retrieve no messages at all because the **Pre-Encryption for HTTPS** scenario focuses on HTTP messages only and the **Network Analysis Parsing Level** stops when parsing of the Network Layer is complete, with the exception of certain other traffic on specified ports.

However, if you are interested in diagnosing the Network Layer, along with certain Transport Layer messages, and Application Layer traffic is of no consequence, then either a **Loopback and Unencrypted IPSEC** or **Local Network Interfaces Trace Scenario** working together with the **Network Analysis Parsing Level** should produce a desirable result while reducing the overhead associated with capturing and parsing Application Layer protocols.

Parsing Level Impact on Analysis Perspectives

After you choose a **Parsing Level**, start a Live Trace Session, and then observe data displaying in the default **Analysis Grid** viewer, you will see a **Parsing Level** indicator on the Message Analyzer Status Bar below the grid. If you are loading saved data through the **Open** feature (on the Message Analyzer **File** menu or the global Message Analyzer toolbar) from a trace file that reflects a **Parsing Level** set by another user, you can quickly ascertain what that **Parsing Level** is by looking at the Status Bar setting after the data is loaded. This information can be important to your analysis perspective, because it can alter the message set from what you might ordinarily expect in a particular **Trace Scenario**. In addition, you should also note that the **Message Stack Tool Window** and the inline origins display (click the green cubes icon on any top-level message in the **Analysis Grid** viewer) will reflect the **Parsing Level** by not showing messages above the top-most parsed level. Note that the resulting simplified view of the message stack can also improve your analysis processes.

Working with Viewpoints

By setting a **Parsing Level**, you can achieve dramatic performance gains, while at the same time you change the way message content displays. By setting a **Viewpoint**, messages are reorganized to show data from the viewpoint of a particular protocol, which also changes the way message content displays. When working with **Viewpoints** and **Parsing Levels**, you should consider the **Parsing Level** of the data when applying a **Viewpoint**. For example, if you configured the **Network Analysis Parsing Level** for a particular set of messages displaying in the **Analysis Grid** viewer and you then set the Application Layer **HTTP** or **SMB/SMB2** viewpoint from the **Viewpoints** drop-down list on the Message Analyzer Filtering toolbar, no results will display in the **Analysis Grid** viewer. This is because the Network Layer is the top-most level to which Message Analyzer parses in the **Network Analysis** scenario, which is well below the indicated Application Layer viewpoints. On the other hand, if you select a **Viewpoint** such as the **TCP Layer**, Message Analyzer will provide a precisely focused view of TCP messages only. Therefore, you might consider the following factors to maximize the effectiveness of using these features together:

- Prior to starting a Live Trace Session, choose a **Parsing Level** that will work with the **Viewpoint** that you have in mind.
- Set the **Parsing Level** up-front to maximize performance gains.
- Set **Viewpoints** to obtain focused analysis in the resulting message set.

Interacting with the Truncated Parsing Mode

If you are loading files that contain truncated messages, such as .cap, .pcap, .pcapng, or .matp files; or if the **Truncated Parsing** option is set when loading data into Message Analyzer, you might encounter some issues in your trace results, depending on the **Parsing Level** that you set. If Message Analyzer detects truncated messages in a file containing data that you want to load, or if you specifically set the **Truncated Parsing** option on the **Files** tab of the **New Session** or **Edit Session** dialog for a Data Retrieval Session, Message Analyzer

automatically switches to a limited OPN parser set consisting of parsers for the Ethernet, GRE, IPv4, IPv6, ESP, AH, IKE, AIPS, TCP, UDP, and HTTP protocols. Truncated parsing can improve performance by reducing the number of messages that Message Analyzer parses and displays; however, you will need to consider the resulting limited message set if you also configure a **Parsing Level**. For example, when the **Truncated Parsing** mode is active and a **Parsing Level** such as **File Sharing** or **Identity and Active Directory** is also set, you will be unable to see SMB or LDAP messages, respectively, in the trace results displayed by a viewer such as the **Analysis Grid**. This is because the limited OPN parser set that is used in the **Truncated Parsing** mode does not include parsers for these protocols.

Therefore, depending on the analysis perspective that you require, it might be best to limit the use of the **Truncated Parsing** and **Parsing Level** features together. Both features provide performance gains, which in either case is a positive outcome, but regarding the impact on focused analysis, you should consider that the advantages of one may outweigh that of the other, depending on the **Parsing Level** that you specify.

Saving and Viewing Trace Data with Set Parsing Levels

If you run a Live Trace Session with a particular **Parsing Level** set, the level displays in the **Parsing Level** label on the Message Analyzer Status Bar below the **Analysis Grid** viewer. You can save your trace results as you normally do in a *.matp file and the saved message set will reflect the applied **Parsing Level**. However, although you can load and view the data from such a file through a Data Retrieval Session, you can view the original **Parsing Level** results that exist in the retrieved data only if you load the file into Message Analyzer through the **Open** feature that is accessible from the Message Analyzer **File** menu or from the global Message Analyzer toolbar. If you load the data from the saved file in a Data Retrieval Session with the **New Session** dialog, by default, Message Analyzer reparses the data with the **Parsing Level** set to **Full**, in which case, the original **Parsing Level** in the saved message set is not reflected in the resulting display of messages.

Obtaining Parser Level Updates

The **Parsing Level** scenarios that you select from the **Parsing Level** drop-down list in the **New Session** or **Edit Session** dialog are Message Analyzer assets. Similar to **Viewpoints** and other asset libraries, you can either download the **Message Analyzer Parsing Level** asset collection once, or auto-sync the collection to receive periodic collection updates that Microsoft provides through a web service. This service integrates with Message Analyzer through the Sharing Infrastructure. You can interact with this infrastructure from the **Asset Manager** dialog, which is accessible from the Message Analyzer **Tools** menu.

More Information
To learn more about the Truncated Parsing mode, see Detecting and Supporting Message Truncation.
To learn more about **Viewpoints**, see Applying and Managing Viewpoints.
To learn more about downloading **Parsing Level** asset collections or auto-syncing to periodic collection updates as part of the Message Analyzer Sharing Infrastructure, see Downloading Assets and Auto-Syncing Updates.

See Also

Specifying a Parsing Level

Specifying Advanced ETW Session Configuration Settings

5 minutes to read

This section describes how to use the **ETW Session - Advanced Configuration** dialog when configuring a Live Trace Session, to enhance the performance of event tracing in the underlying ETW Sessions that all Message Analyzer providers utilize. To facilitate this enhancement, the **ETW Session - Advanced Configuration** dialog enables you to modify configuration settings for the ETW Session in which the ETW Provider components of a Message Analyzer Live Trace Session participate. Mainly, this involves adjusting settings for the ETW buffer configuration of the ETW Session that is managed by the ETW Controller. You can open the **ETW Session - Advanced Configuration** dialog by clicking the **Configure ETW Session** button on the **Live Trace** tab of the **New Session** or **Edit Session** dialog (if you are editing an existing session). The **Configure ETW Session** dialog is shown in the figure that follows.

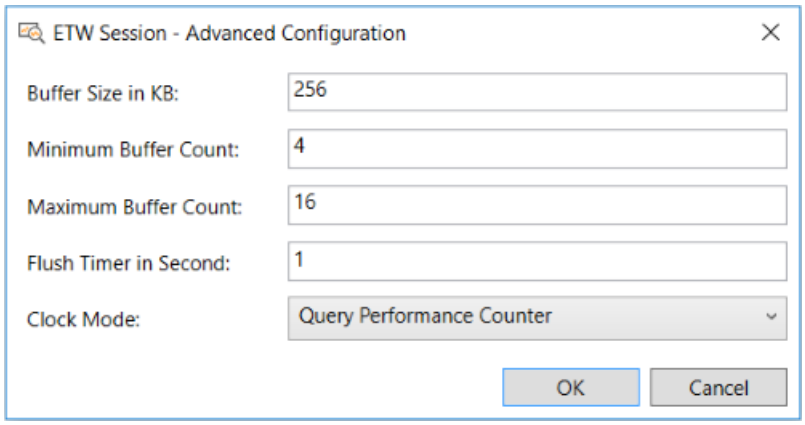

Figure 28: ETW Session - Advanced Configuration

To set the context for using the **ETW Session - Advanced Configuration** dialog, some background information about ETW buffers is included here.

> **NOTE**
>
> The most common reason why you might want to specify **ETW Session - Advanced Configuration** settings is that your Live Trace Sessions are dropping packets, the evidence of which could be spurious diagnosis errors, excessive TCP retransmits or broken three-way handshakes, dropped packet log indications for a **Microsoft-PEF-WFP-MessageProvider** trace, and other occurrences where trace results seem odd.

ETW Buffer Background Concepts

The **ETW Session - Advanced Configuration** dialog enables you to specify values for key ETW buffer settings. By adjusting these settings for an ETW Session, you can reduce the risk of dropping event packets during a Live Trace Session with Message Analyzer. A properly configured ETW Session will prevent the loss of event data, although acquiring the right configuration can be tricky. For example, larger ETW buffers are preferred over smaller buffers to enable more efficient disk I/O as buffer data is written to disk. However, as you specify higher values for the ETW buffer settings more memory will be consumed, thus impacting performance.

Events are captured in an ETW Session by accepting them from the ETW Provider API. An ETW Session is a collection of in-memory buffers that are managed by the kernel. At all times, an *in-use* buffer is assigned to each processor. Each time the EventWrite() method is called, space is reserved in the in-use buffer that is currently

allocated to the processor on which the calling thread is running. The event header and user data is then copied into that space. When the in-use buffer becomes full, it is flushed to the ETW Session's log file and/or to the ETW Consumer in real time. At this point, a free buffer from the buffer pool is assigned to the processor. If the logging throughput exceeds the ability of the flusher to free up buffers—for example, because the incoming event throughput is greater than the disk write throughput—at some point all available buffer space in the ETW Session might be consumed, causing the EventWrite() method to fail with an ERROR_NOT_ENOUGH_MEMORY message along with a loss of event data.

To avoid this loss of event data, it is critical that the flusher interval is adequate, enough ETW buffers are available from the buffer pool, and that buffer size is set correctly to ensure good I/O write efficiency.

Configuring the ETW Session

When you are configuring a Live Trace Session, you have the option to specify several settings in the **ETW Session - Advanced Configuration** dialog. These settings provide information to the ETW Session Controller that enable it to control various aspects of an ETW Session to ensure that data is not lost from dropped event packets. After you create a new Live Trace Session and you have selected a **Trace Scenario** in the **New Session** dialog, click the **Configure ETW Session** button in the **New Session** dialog to display the **ETW Session – Advanced Configuration** dialog. You can then configure values for the following settings:

- **Buffer size in KB** — sets the size of all ETW buffers in the ETW buffer pool, to which an **ETW Provider** writes events during an ETW Session.

  Buffer size has an impact on disk I/O write efficiency. For example, small ETW buffer sizes can reduce I/O write efficiency. However, the Message Analyzer default setting of 256 KB for this value should ensure good write performance, reduction in disk overhead, and the decreased likelihood that events will be lost. Buffer size also determines the maximum size of any event that can be traced. It is typically unexpected that a single event will be so large that it cannot fit into an ETW buffer. However, if buffer size is set too low and an event is unusually large, it may never be traced. In the ETW framework, event size is limited to slightly less than 64 KB, so if you set buffer size equal to or greater than 128 KB, no events will be lost.

- **Minimum Buffer Count** — specifies the minimum number of buffers to use when writing the events of a provider.

- **Maximum Buffer Count** — specifies the maximum number of buffers in the buffer pool that are available for writing the events of a provider.

  Buffer count is a key factor in determining the capacity of an ETW Session. If you set the buffer count too high, memory will be wasted, whereas if you set the buffer count too low, a large number of events could be lost. The **Minimum Buffer Count** represents the smallest possible footprint of an ETW Session. Although it can be a process of trial and error, getting this number right is critical because a small buffer pool can be rapidly consumed as events are written, especially if the buffer size is also small, and this can lead to a loss of data.

- **Flush Timer in Seconds** — specifies the interval at which in-use ETW buffers allocated from the buffer pool are to be flushed. The critical issue here is to ensure that the event throughput does not exceed the ability of the flusher to free up ETW buffers. The **Flush Timer** interval should coincide with the moment that the in-use ETW buffer is full, so it can be flushed and its data can be sent to the ETW Consumer (Message Analyzer in this case) so that a new ETW buffer can be assigned to the session and event tracing can continue without the loss of data.

- **Clock Mode** — enables you to specify different time resolution values for the event timestamps of an ETW Session, based on the following clock types:

  - **Query performance counter** — the default clock type, which can have a resolution of 100 nanoseconds or better.

- **System time** — typically around 10 milliseconds resolution.
- **CPU tick count** — time resolution is based on CPU cycles.

Decrypting TLS and SSL Encrypted Data

12 minutes to read

In addition to the many tools that Message Analyzer provides to filter, analyze, and visualize network traffic and other data, Message Analyzer also provides a **Decryption** feature that can help you diagnose traces that contain encrypted Transport Layer Security (TLS) and Secure Sockets Layer (SSL) traffic. Decrypting TLS/SSL traffic can be critical to troubleshooting network, protocol, performance, and connectivity issues. The Message Analyzer **Decryption** feature also resolves existing limitations of the **Microsoft-PEF-WebProxy** Fiddler message provider, such as the non-transparency of errors and the inability to capture other TLS/SSL encrypted traffic besides HTTPS.

> **NOTE**
>
> Please be aware that decryption of TLS v1.2 messages is also supported, as is TDS/TLS.

The Message Analyzer **Decryption** feature enables you to view data for Application layer protocols that are encrypted with TLS and SSL, such as the HTTP and Remote Desktop (RDP) protocols. However, to enable a **Decryption** session in Message Analyzer, you will need to import a certificate that contains a matching identity for a target server, specify a required password, and then save the configuration. You can then either load a saved trace file into Message Analyzer through a Data Retrieval Session or start a Live Trace Session that will be enabled for decryption. Thereafter, Message Analyzer decrypts the trace by using the server certificate and password that you provided. After the trace results display in the **Analysis Grid** viewer, a **Decryption Tool Window** holds the decryption analysis information. If there are decryption failures, errors are reported to the **Decryption** window, where a red **Error** icon displays for each message that failed the decryption process. Detailed error descriptions are also provided in the **Decryption** window to assist in troubleshooting and analysis. If there are no errors reported, then the **Decryption** window displays either a blue **Info** icon for each message that was successfully decrypted, or a yellow **Warning** icon that flags each message for which a certificate could not be found.

Workflow Overview

Decryption works within the existing Message Analyzer architecture to simplify its usage and results analysis. After you provide and save one or more server certificates and passwords, Message Analyzer will decrypt target traffic that is encrypted with the Transport Layer Security (TLS) or Secure Sockets Layer (SSL) security protocols for any session containing such traffic in the current Message Analyzer instance. Note that existing certificates are recovered from the certificate store after Message Analyzer restarts; however, for security purposes, passwords are not. Therefore, if you want decryption to occur after a Message Analyzer restart, you will need to manually re-enter passwords each time. In any single instance of Message Analyzer where you have entered passwords for existing or new certificates, Message Analyzer can decrypt target messages whenever you load a saved file or run a live trace that retrieves such messages. Thereafter, Message Analyzer displays the decrypted data in the **Analysis Grid** viewer at an upper layer of the protocol stack and provides a separate **Decryption** window that presents decryption session analysis information, including status and errors. You can also save a trace that contains decrypted data in the .matp format, just as you would any other trace. All the Message Analyzer tools and features that normally enable you to manipulate and analyze message data are available for use in a decryption session, including the **Details**, **Message Data**, **Field Data**, and **Message Stack Tool Windows** that enable you to focus on specific message fields, properties, values, and layers.

The following steps are an overview of the workflow that you will typically follow when working with the **Decryption** feature:

1. Import and store server certificates and add passwords as required on the **Decryption** tab of the **Options** dialog that is accessible from the Message Analyzer global **Tools** menu, as described in Adding Certificates and Passwords.
2. Start a Live Trace Session or load a saved file through a Data Retrieval Session that contains target messages to enable Message Analyzer to decrypt as many conversations as possible, as described in Decrypting Trace Data.
3. View decryption status information and analyze results in the **Decryption** window grid, as described in Analyzing Decryption Session Data.
4. Select message rows in the **Decryption** window and observe corresponding selection of decrypted messages in the **Analysis Grid**, as described in Viewing Decrypted Messages.
5. Use the **Analysis Grid** viewer along with the **Details**, **Message Data**, and **Message Stack** windows to analyze the decrypted data, as described in Viewing Decrypted Messages.
6. Save a decrypted trace in .matp format for sharing with others or for use in other applications, as described in Saving Decryption Session Data.

Adding Certificates and Passwords

To add a server certificate to the Message Analyzer certificate store, you must add it to the grid of the **Certificates** pane on the **Decryption** tab of the **Options** dialog that is accessible from the Message Analyzer **Tools** menu. To import a certificate into the Message Analyzer certificate store, click the **Add Certificate** button on the toolbar of the **Decryption** tab to open the **Add Certificate** dialog, navigate to the directory where the certificate is located, select the certificate, and click the **Open** button to exit the dialog. Each time you add a certificate in this manner, the **Password** field displays to enable you to manually enter a password for the certificate you are adding. The **Decryption** feature can decrypt conversations only if a corresponding certificate exists in the store and a password is provided for it. If you do not enter a password, or if it is an incorrect password, you will be prompted to add the correct information.

All certificates and passwords that you add to the grid on the **Decryption** tab of the **Options** dialog are saved to the certificate store and persist in the current Message Analyzer instance, unless you remove them by clicking the **Clear List** button in the toolbar on the **Decryption** tab. Note that if you remove certificate entries from the list and click **OK** to exit the **Options** dialog, neither the certificates nor the passwords will be listed in the grid following a Message Analyzer restart.

Obtaining a Server Certificate

If you require a security certificate from a server on which you are capturing encrypted message traffic, you can obtain one by using the Certificate Manager MMC to export a server-side certificate (from the server). Note that the client-side version of such a certificate does not have the information needed to decrypt the data. Cipher suites are typically decided by the server configuration. Note that you may have to make modifications to the client- or server-side to locate a cipher suite that is supported by Message Analyzer.

Enabling Certificates

Certificates and passwords are not enabled for a Data Retrieval Session or a Live Trace Session unless you specifically select them prior to the session. To enable or disable a certificate in the certificate list, either select or unselect the check box to the left of the certificate name, respectively. In addition, you can enable or disable all certificates simultaneously by either selecting or unselecting the **Name** check box, respectively. Before loading data or running a live trace that you want to decrypt, make sure that your certificate/s are enabled by selecting the appropriate certificate check box in the certificate list; you should also ascertain that a password for the certificate displays in hidden text in the **Password** column. When you are finished adding certificates and

passwords, click the **OK** button to exit the **Options** dialog, at which time the certificate and password are stored. Thereafter, the certificate that you added is accessible to all subsequent decryption sessions, including those following Message Analyzer restarts; however, the password is made available for decryption sessions that use that certificate in the current Message Analyzer instance only.

Decrypting Trace Data

After adding the appropriate certificate/s to the grid on the **Decryption** tab of the **Options** dialog and specifying the corresponding password/s, you are ready to load data into Message Analyzer or to run a live trace that retrieves the applicable messages that you want to decrypt. You can load data through a Data Retrieval Session or run a Live Trace Session with decryption enabled in the same way you load any other file or run any other trace with Message Analyzer. Prior to loading a file in a Data Retrieval Session or starting a Live Trace Session with decryption enabled, you are advised to select the **Analysis Grid** viewer from the **Start With** drop-down list in the **New Session** dialog for ease of analysis. After you load your data or you stop a Live Trace Session and the **Decryption** process completes, message data displays in the **Analysis Grid** viewer and the decryption analysis data displays in the **Decryption Tool Window**.

Message Analyzer also makes it convenient for you to decrypt messages that were captured in multiple simultaneous conversations across different servers. Rather than decrypting the conversations one by one, you can simply input a list of certificates and passwords into the certificate store for the current Message Analyzer instance and then load the trace files containing the target conversations into Message Analyzer through a Data Retrieval Session. Thereafter, Message Analyzer will match the provided certificates to the appropriate conversations and decrypt as many messages as possible.

Analyzing Decryption Session Data

After Message Analyzer decrypts messages in a trace, you can analyze the results in the **Decryption** window. If this **Tool Window** is not already open, you can display it by clicking the Message Analyzer **Tools** menu, selecting the **Windows** item, and then clicking **Decryption** in the menu that displays. The **Decryption** window contains the following columns of information:

- **Type** — lists the different **Decryption** types with different icons, which have the meanings that follow. Note that messages that are associated with these icons are displayed in the **Message** column:
  - **Info** — this blue icon indicates that a message was successfully decrypted, as the necessary certificates and passwords were located in the certificate store. The informational message associated with this icon might consist of the name of the certificate that was used in the decryption.
  - **Warning** — this yellow icon indicates that a message could not be decrypted, as the corresponding certificate and password were unavailable. The warning message associated with this icon might consist of public key information for the missing certificate that can help you to locate it.
  - **Error** — this red icon indicates that a message could not be decrypted because an error occurred, for example, as the result of an incomplete TCP handshake. The error message associated with this icon might consist of error information and troubleshooting advice.
- **Timestamp** — specifies message time stamps in a user friendly format.
- **ProtocolVersion** — lists the version of the protocol message that was subject to the **Decryption** process.
- **SourcePort** — displays the TCP source ports that carried the conversations that were subject to the **Decryption** process.

- **DestinationPort** — displays the TCP destination ports that carried the conversations that were subject to the **Decryption** process.
- **SourceAddress** — displays the source IP addresses that carried the conversations that were subject to the **Decryption** process.
- **DestinationAddress** — displays the destination IP addresses that carried the conversations that were subject to the **Decryption** process.
- **Decrypted Message Count** — lists the number of messages that were decrypted.
- **Undecrypted Message Count** — lists the number of messages that were not decrypted.
- **Message** — specifies informational messages related to one of the decryption icons, to provide additional details about the **Decryption** process, for example **Error** information.

Viewing Decrypted Messages

To view the decrypted data itself, select a message row in the **Decryption** window and observe that it drives message selection in the **Analysis Grid** viewer. For example, if Message Analyzer decrypted HTTPS data, a top-level HTTP message might now be selected in the **Analysis Grid** viewer, as shown in the figure that follows, depending on the message row you selected in the **Decryption** window. Application Layer HTTPS messages are normally encrypted when the Microsoft Windows Internet (WinInet) Service passes them to cryptographic components such as the TLS security protocol, which in turn passes encrypted/secure data to the Transport Layer. As a result of TLS encryption, Message Analyzer does not normally process and display such messages unless you use the **Decryption** feature. When you do, you can see the encrypted message layer as part of the origins stack beneath the decrypted message at top-level in the **Analysis Grid** viewer. In the example of the following figure, the node of the decrypted HTTP message is expanded so that you can see the encrypted message layer (TLS) and the original protocol stack in the origins tree.

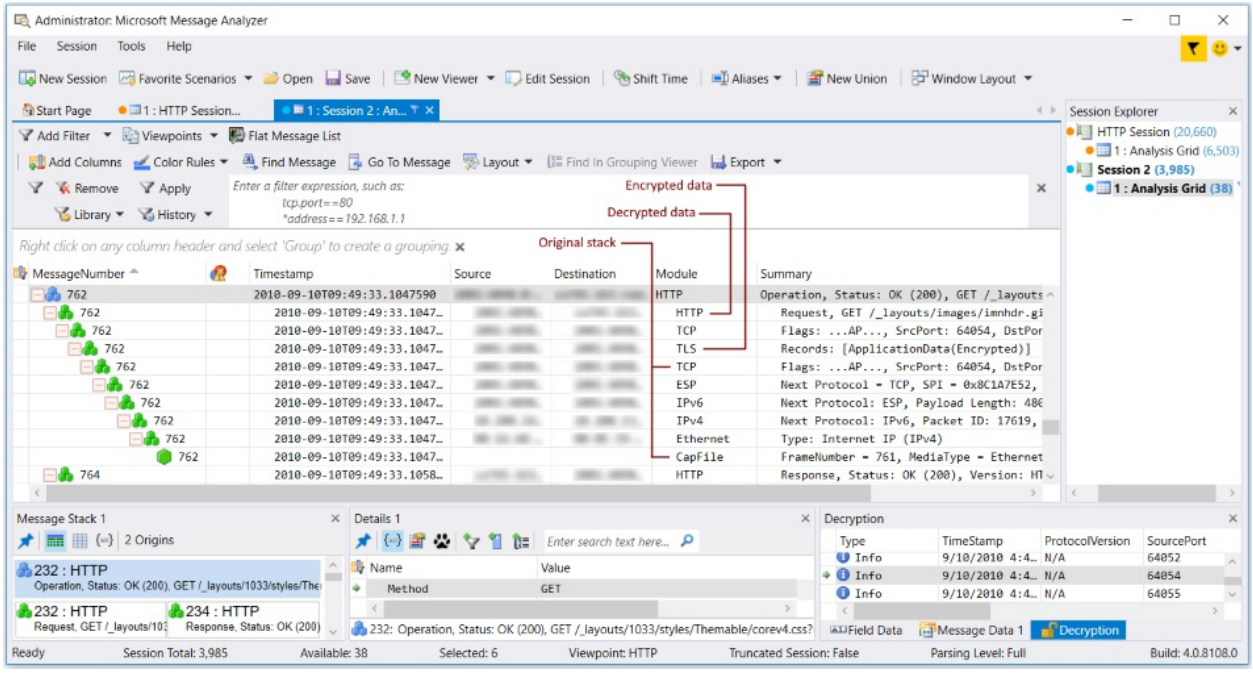

Figure 29: Decrypted data and message stack

Because message selection in the **Analysis Grid** viewer in turn drives the display of data in the **Details**, **Message Data**, and **Message Stack** window, you can use these **Tool Windows** to view the decrypted and encrypted message data, which includes field data in the **Details**, hexadecimal data in the **Message Data**, and top-level messages with origins layers in the **Message Stack** window. Note that field selection in the **Details** window drives the display of corresponding hexadecimal data in the **Message Data** window, such that you can view the hexadecimal value of any message field that you select.

Saving Decryption Session Data

You can save decrypted trace data in the same manner that you save any other trace data. You simply select the **Save** item in the Message Analyzer **File** menu to display the **Save As** dialog, where you can select one of three save options, as follows:

- **All Messages** — enables you to save all messages from a **Decryption** session that are displayed in the **Analysis Grid** viewer.
- **Filtered Messages** — enables you to save a filtered message set, for example, if you applied a view **Filter** to your **Decryption** session results.
- **Selected Messages** — enables you to save only specific messages that you select in an **Analysis Grid** viewer session tab; for example, you might want to save specific decrypted messages and their origins messages of the associated protocol stacks.

After you specify one of the indicated save options, you can save your data as follows:

- **Save As** — click this button to save your decrypted data in the native .matp trace file format. If you reload such a file, it is unnecessary to enable **Decryption** again to see the decrypted data in a Message Analyzer viewer such as the **Analysis Grid**.
- **Export** — click this button to save your data in the .cap trace file format. If you select this option, your decrypted data will not be saved; however, you can reapply **Decryption** to this file when you reload it so that you can view the decrypted data again.

More Information
To learn more about starting and configuring a new Data Retrieval Session or a Live Trace Session, see Starting a Message Analyzer Session.
To learn more about saving trace data, see the Saving Message Data section.

See Also

Decryption Tool Window

Selecting a Session Data Viewer

4 minutes to read

Message Analyzer provides a set of native data viewers for all Message Analyzer installations. The features and functions of these data viewers are described in detail in the Data Viewers topic. This section briefly describes the viewers that are available, selection locations, selection criteria, and how to identify multiple viewers that are common to a particular session.

Choosing a Session Viewer

Before you start a Live Trace Session or Data Retrieval Session, you can select the data viewer with which to view your trace results data. These data viewers are available from the **Start With** drop-down list that appears in the **New Session** dialog that you open when configuring a Live Trace Session or Data Retrieval Session. The data viewers that you can select from this location consist of the following:

- **Analysis Grid**
- **Grouping**
- **Pattern Match**
- **Gantt**
- **Chart** — displays a default **Bar** element graph only when chosen from this location.
- **Interaction**
- **Message Summary Tiles***
- **Message Summary Lists***
- **Perfmon***

Note that the data viewers in this list that are marked with an asterisk (*) are **Preview** features. To make these viewers available for selection during session configuration, you must select them on the **Features** tab of the **Options** dialog, which is accessible from the Message Analyzer **Tools** menu, and then you will need to restart Message Analyzer.

> **NOTE**
> If you do not specify a data viewer prior to starting a Live Trace Session or Data Retrieval Session, then Message Analyzer uses the current default viewer to display data, for example, the **Analysis Grid** viewer. You can set the default viewer for all session results in the **Default Profiles** pane on the **Profiles** tab of the **Options** dialog.

Other Viewer Selection Locations

You can also select the previously indicated viewers from either of the following alternate locations to change the data view for a set of trace results during an Analysis Session:

- **New Viewer** drop-down list — after you start a Live Trace Session or Data Retrieval Session, you can specify a different viewer for your data by clicking the **New Viewer** item in the global Message Analyzer **Session** menu or by clicking the **New Viewer** button on the global Message Analyzer toolbar to display a drop-down list that contains all the viewer selections.

- **Session Explorer Tool Window** context menu — after you start a Live Trace Session or Data Retrieval Session, you can specify a different viewer for your data by right-clicking anywhere in the **Session Explorer** window, clicking the **New Viewer** item in the context menu that appears, and then selecting a data viewer of choice along with a **Layout** if appropriate.

Selecting Layouts
When you select data viewers from these alternate locations, the viewers listed immediately below will also enable you to choose a built-in view **Layout** in which to present the data. The built-in **Layouts** are custom designed by Microsoft to expose data that is relevant to specific types of analysis. For example, if you are interested in analyzing file sharing performance data after performing a capture, you might choose the **File Sharing Perf SMB/SMB2 Layout** for the **Analysis Grid** viewer in the **New Viewer** drop-down list on the global Message Analyzer toolbar.

- **Analysis Grid**
- **Grouping**
- **Chart**
- **Charts (Deprecated)**

Selection Criteria

Depending on what you wish to achieve, there may be some things to consider when selecting a viewer in which to display message data that you captured live or retrieved from saved files. For example, the selected viewer should to some degree reflect the type of data being captured and should be considered in terms of the problem/s you are trying to isolate.

For example, if you simply want to expose data in a standard analysis environment, you can display data in the default **Analysis Grid** viewer **Layout**, as described in the Analysis Grid Viewer topic. If you want to analyze process IDs and the associated IP conversations, you could choose the **Network Conversation Tree with Process ID Layout** for the **Analysis Grid** viewer, as described in Applying and Managing Analysis Grid Viewer Layouts. Moreover, you can make similar choices for the **Grouping** viewer and the **Chart** viewer, which are each accessible from the same **New Viewer** drop-down list during analysis of session results.

If you want to look at high-level data summaries and statistics, you might consider using one of the many **Layouts** for the **Chart** viewer, for example, the **IP/Ethernet Conversations by Message Count Top 20** layout that provides a graphic **Bar** element representation of the relative distribution of message volume for each IP conversation in a set of trace results. If you suspect network connection or latency issues, you might want to open the **TCP Diagnosis with Stevens** and/or **Top TCP/UDP Conversations Layouts** for the **Chart** viewer.

Identifying Session Viewer Data

Whenever you select a data viewer for a *new* session, it displays in a new and separate session viewer tab with a unique name just below the global Message Analyzer toolbar. If you select a new data viewer for an *existing* session, it displays in a new session viewer tab with the same name as the existing session, also just below the global Message Analyzer toolbar. To maintain correlation between data viewers of the same session, Message Analyzer provides an identical colored dot on the session tab of each data viewer in the same session, for ease of identification.

More Information
To learn more about the different types of data viewers that are available, along with the **Tool Windows** with which they interact, see the Data Viewers and Tool Windows topics of this Operating Guide.
To learn more about creating custom **Layouts** for the **Chart** viewer, see Extending Message Analyzer Data Viewing Capabilities.

See Also

Selecting a Data Retrieval Session Viewer
Session Data Viewer Options

Creating Remote Session Configurations

2 minutes to read

This section describes how to set up a **Trace Scenario** that will capture data on a target remote computer. It begins with operating system requirements and then provides instructions on how to configure the source and host computers for a remote capture, which includes configuring the WinRM service and Trusted Hosts list. Next, you are shown how to specify one or more target computers for remote capture, along with how to specify host connection data.

To provide a data source for your remote capture session, you will need to specify a remote data provider, which in this case is the **Microsoft-Windows-NDIS-PacketCapture** provider that uses Windows Management Instrumentation (WMI) for its remote capabilities. By default, this provider is contained in the built-in **Remote Network Interfaces Trace Scenario** that you can select from the **Select Scenario** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog during Live Trace Session configuration. After you select this **Trace Scenario**, you can specify advanced settings for the **Microsoft-Windows-NDIS-PacketCapture** provider from the associated **Advanced Settings** dialog, as indicated below. This includes enabling and disabling adapters on which to capture, optionally specifying promiscuous mode for supporting adapters, specifying various types of low-level stack or Hyper-V-Switch extension layer filters, packet direction filters, and configuring various types of special filters such as **EtherTypes** and **IP Protocol Numbers**.

The content is covered in the following topics, which includes detailed descriptions of how to use the **Advanced Settings** dialog for the **Microsoft-Windows-NDIS-PacketCapture** provider:

Configuring a Remote Capture
Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog

See Also

Microsoft-Windows-NDIS-PacketCapture Provider
Configuring a Live Trace Session

Configuring a Remote Capture

10 minutes to read

Message Analyzer enables you to capture traffic on remote machines by running the **Remote Network Interfaces Trace Scenario**. This scenario uses the **Microsoft-Windows-NDIS-PacketCapture** provider for remote tracing. However, you can capture remote traffic with the **Remote Network Interfaces** scenario only from computers that are running the Windows 8.1, Windows Server 2012 R2, Windows 10, or later operating system. In addition, the source machine from where you start the **Remote Network Interfaces** scenario must also be running one of these same operating systems. Note that the **Remote Network Interface** scenario is not available on computers that are running the Windows 7, Windows 8, or Windows Server 2012 operating system.

Other requirements also apply, as follows:

- **WinRM configuration** — this service requires configuration on both the source computer where you are running the Message Analyzer remote trace and on target computers from which you are capturing data. You can configure the WinRM service by running the following command string from an elevated command prompt (Run as Administrator):

  ```
  winrm quickconfig
  ```

- **Trusted Hosts configuration** — when the source computer and remote target host/s are not in the same domain, you must add the remote host name to the source computer Trusted Hosts list by running the following command string from an elevated command prompt, while substituting appropriately for the *RemoteHostName* value:

  ```
  winrm set winrm/config/client @{TrustedHosts="RemoteHostName"}
  ```

  > **IMPORTANT**
  > If you intend to capture data from multiple remote hosts, you should use the previous command to specify them in comma-separated format surrounded in quotes, as follows:
  > `winrm set winrm/config/client @{TrustedHosts="RemoteHost1Name, RemoteHost2Name,.."}`, and so on

Connecting with Remote Hosts

If your systems match the indicated requirements, you can start remote configuration for your Live Trace Session by first specifying the name of one or more remote Windows 8.1, Windows Server 2012 R2, or Windows 10 hosts on which to capture remote message traffic, possibly along with Administrator credentials (see Specifying Remote Host Connection Data), to connect with those computers. You can specify the connection information from the **Edit Target Computers** dialog, which is accessible by clicking the **Edit** button on the **Live Trace** tab of the **New Session** dialog. If you do not specify any remote computers on which to capture traffic, then your Live Trace Session defaults to capturing messages on the local computer, as indicated by the default **Localhost** setting in the **Target Computers** list on the **Live Trace** tab.

Moreover, Message Analyzer enables you to capture traffic with the following target configurations:

- **Localhost** — the default setting to capture message traffic on the local computer.
- **One or more remote computers** — remove the **Localhost** setting and add target remote computers on which to capture message traffic concurrently to the **Target Computers** list.

- **Local and remote computers** — retain the **Localhost** setting and add remote computers on which to capture message traffic concurrently to the **Target Computers** list.

> **NOTE**
>
> When you capture messages from multiple computers, the results are aggregated in the data viewer that you specified in the **Start With** drop-down menu of the **New Session** dialog. If this is the **Analysis Grid** viewer, you can add the **DataSource** field to your trace results display from the **General** node of **Field Chooser** (right-click **DataSource** and select **Add as Grouping**) to group and therefore differentiate the messages from each data source (remote computer).

Specifying Remote Host Connection Data

To provide remote host connection information, click the **Add** drop-down in the **Edit Target Computers** dialog and then select **New Row** from the menu. A new row is then added to the **Edit Target Computers** dialog for specifying the host name or IP address in the **Computer Name / IP Address** text box, along with required connection credentials in the **User Name** and **Password** text boxes. If you are connecting to multiple remote hosts and the connection credentials are identical, you should select the **Automatically copy credentials** check box to avoid repeating duplicate credential entries for each host.

If you can use your current logon credentials to connect with specified hosts, leave the **User Name** and **Password** text boxes blank. You can do this if you will be connecting to a remote computer in the same domain and your local logon credentials have sufficient access (Administrative) privileges to the remote computer. Otherwise, you should specify your user name in the Domain\Username format and provide a valid password. You will still need sufficient access privileges in this case to connect with the remote computer that is in a different domain. When you are done adding host connection information, click **OK** to exit the **Edit Target Computers** dialog.

If you want to edit the connection information for any of the hosts that display in the **Target Computers** list, simply click **Edit** on the **Live Trace** tab of the **New Session** dialog and modify the information as required. Note that Message Analyzer retains all the connection information that you specify from session to session, until you remove it with the **Delete** control in the **Edit Target Computers** dialog. In addition, all the host names and user names that you entered in the **Edit Target Computers** dialog are retained in the **Add** drop-down list for future selection as required.

Specifying Advanced Settings for Remote Hosts

After you have specified connection credentials for remote computers on which to capture message traffic, and you have selected the **Remote Network Interfaces Trace Scenario**, you have the option to specify special filtering configurations in the **Advanced Settings – Microsoft-Windows-NDIS-PacketCapture** dialog. To access this dialog, click the **Configure** link to the right of the **Microsoft-Windows-NDIS-PacketCapture** provider **Id** in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog. When the dialog opens, select the **Provider** tab and then click the **Host** drop-down to display the list of target computers that your Live Trace Session will connect with. To specify different filtering configurations for different hosts, you must select each host separately in the **Host** drop-down and then specify the filters you want to use, as described in Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

Enumerating Adapters

When you select an item in the **Host** drop-down list, Message Analyzer attempts to connect with the specified host computer that is running any of the following operation systems: Windows 8.1, Windows Server 2012 R2, or Windows 10. If the connection attempt is successful, Message Analyzer may display a progress bar while enumerating the network adapters on the selected remote host, which can include both host adapters and virtual machine (VM) adapters that are serviced by a Hyper-V-Switch. Message Analyzer returns the results of the enumeration to the Interface Selection grid of the **Advanced Settings** dialog and populates it

with adapter **Name** and **MAC Address** information. Thereafter, when you open the **Advanced Settings** dialog in the current session, the tree grid section of the dialog is automatically populated with this information. You can then apply the filtering configurations that you specify to one or more enumerated adapters that you select in the **Advanced Settings** dialog.

If you want to specify unique filtering configurations for each remote host, you will need to repeat the previously indicated process for each host in the **Host** list on the **Provider** tab of the **Advanced Settings** dialog. However, to simplify the process, you have the option to apply a common filtering configuration to all hosts by clicking the **Apply Changes to All Hosts** button on the **Provider** tab. Before you do this, you should carefully consider the applicability of the filtering configuration that you apply to all adapters on the remote hosts. For example, certain filters have a different effect on the interception layer and packet traversal path through the NDIS stack for host adapters, versus the extension layers of a Hyper-V-Switch adapter that services one or more virtual machines (VMs). For more details on the effects of **Layer** filtering and packet traversal **Direction** filters, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

Capturing Remotely in Promiscuous Mode

If you want to capture data remotely with an adapter that supports Promiscuous Mode, you can use the **Advanced Settings — Microsoft-Windows-NDIS-PacketCapture** dialog to do so. In the Interface Selection section of the dialog, you can simply select a P-Mode enabled adapter to capture in Promiscuous Mode, as indicated in Selecting Adapter Interfaces.

> **TIP**
>
> Please be aware that you can also capture data on the local computer in P-Mode, but you must use a **Trace Scenario** that contains the **Microsoft-Windows-NDIS-PacketCapture** provider, otherwise, you will not be presented with the **Advanced Settings** dialog for this provider where you can select an adapter that supports P-Mode. You can display this dialog by clicking the **Configure** link that appears next to the **Microsoft-Windows-NDIS-PacketCapture** provider in the **ETW Providers** list of the **New Session** dialog during Live Trace Session configuration,
>
> You can also specify the **Microsoft-Windows-NDIS-PacketCapture** provider separately, outside of a built-in Message Analyzer **Trace Scenario** that contains this provider, by locating it in the **Add System Providers** dialog that displays when you click the **Add Providers** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog.
>
> In either case, you will need to determine an adapter that supports P-Mode and then select that adapter in the **P-Mode** column of the Interface Selection section of the **Advanced Settings — Microsoft-Windows-NDIS-PacketCapture** dialog to enable a Promiscuous Mode capture with the selected adapter.

Filtering and Adapter Configurations

Message Analyzer enables you to specify the remote adapters on which to capture traffic by selecting specific adapters in the tree grid section of the **Advanced Settings** dialog. In addition, you can specify various filters such as **All Layers**, **Direction**, **Truncation**, **EtherType**, **IP Protocol Numbers**, **IP Addresses**, and **MAC Addresses**.

> **IMPORTANT**
>
> If you want to capture traffic from a target remote VM, you will need to select the VM in the tree grid section of the **Advanced Settings** dialog and then create a **MAC Address** filter based on the VM's MAC address to isolate the data. Otherwise, you will capture Hyper-V-Switch traffic that is destined for all VMs that are serviced by the switch, given that a Hyper-V-Switch driver cannot of itself distinguish between VMs.

Some filters can have a different effect, depending on where you apply them. For example, selecting the **All**

Layers and **Direction** check boxes together enables you to define the layer of the NDIS stack on which packets are intercepted and the traversal path (direction) in which they are intercepted. This can help you determine whether an adapter is dropping packets at a particular NDIS layer. When applied to a switch adapter, you can also define the interception layer and path of packets as they traverse the Hyper-V-Switch Extension layers. However, note that these filters have a slightly different effect, depending on whether you apply them to the NDIS layers of a host adapter versus the Extension layers of a switch adapter. The differences are explained in the **Layer** and **Direction** bullet points of the topic Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

By choosing adapters and setting filters in the **Advanced Settings** dialog, you can be very specific about where you capture remote traffic and how much data you capture. For example, in addition to the previously indicated filters, you could also capture only the packet headers for a particular protocol of interest, by setting a truncation value that matches the header length of that protocol. You can access the configuration of all of the special filters for the **Microsoft-Windows-NDIS-PacketCapture** provider in the **Filters** pane of the **Advanced Settings** dialog, the details of which are further described in Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

> **TIP**
> The **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog configuration settings are tied to the **Microsoft-Windows-NDIS-PacketCapture** provider. As a result, you can use this dialog to specify configuration settings for any **Trace Scenario** that employs the **Microsoft-Windows-NDIS-PacketCapture** provider. You simply access the dialog in the previously indicated manner for this provider. Note that the remote capabilities of the **Microsoft-Windows-NDIS-PacketCapture** provider are available in all **Trace Scenarios** that use this provider on computers that are running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system; but you can only capture remote data successfully if you are targeting computers that are running one of the same operating systems.

Starting a Remote Trace

After Message Analyzer has connected to one or more remote hosts and you have optionally specified adapter and filtering configurations in the **Advanced Settings** dialog, you can close the dialog and start your remote Live Trace Session the same way you start any trace — by clicking the **Start** button in the **New Session** dialog. At this time, Message Analyzer starts multiple concurrent subsessions, that is, if you have successfully connected with multiple hosts for remote capture, where each subsession captures message traffic on a different host that is specified in the **Target Computers** list on the **Live Trace** tab of the **New Session** dialog. The captured data then begins to aggregate from all the subsessions into the **Analysis Grid** viewer, or whichever viewer you specified from the **Start With** drop-down in the **New Session** dialog prior to starting the trace.

See Also

Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog

Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog

12 minutes to read

As described in Configuring a Remote Capture, you can specify the host adapters or VM adapters on which to capture remote traffic, along with various filters and other settings, from the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog. You can open the **Advanced Settings** dialog by clicking the **Configure** link to the right of the **Microsoft-Windows-NDIS-PacketCapture** provider **Id** in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog, after selecting any **Trace Scenario** that uses the **Microsoft-Windows-NDIS-PacketCapture** provider, for example, the **Remote Network Interfaces Trace Scenario**. The **Advanced Settings** dialog consists of selection and configuration settings that are organized into an upper Interface Selection section and a lower Filters section, as shown in the figure that follows.

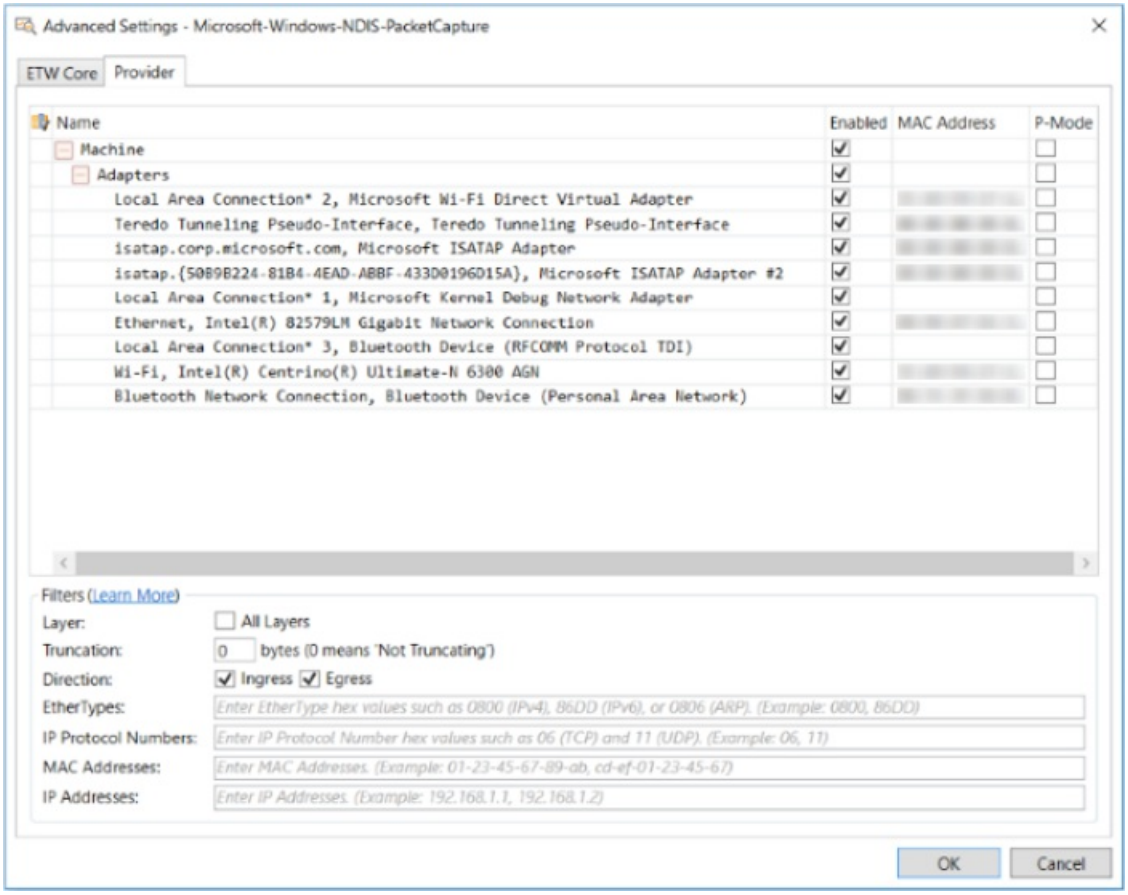

Figure 30: Advanced Settings for the Windows-NDIS-PacketCapture Provider

The configuration components in each of the shown sections are described in this topic.

Selecting Adapter Interfaces

The Interface Selection section of the **Advanced Settings** dialog contains a tree grid with four columns where virtual machine (VM) adapter and/or host adapter information can be populated when you open the dialog. VM adapters are available to virtual machines and host adapters are available to the management operating system (OS). The enumeration of such adapters takes place when you connect to a remote host; thereafter, the tree grid indicates these elements under a collapsible **Machine** node. The four columns in the Interface Selection section consist of the following:

- **Name** — this column contains the names of all adapters that are discovered on the remote host; these names appear under the **Machine** node in this column.
- **Enable** — provides a check box that enables you to capture data on any individual selected adapter or group of adapters. If you select the **Machine** check box in this column, all child elements in the tree grid are automatically selected. If you unselect the **Machine** check box, all child elements in the tree grid are deselected to enable you to choose individual adapters or groups of adapters from which to collect data. For example, you could select the **Switch** node or **Adapters** node to automatically select all the child elements under either of these nodes, to collect data from these sources.
- **MAC Address** — specifies the Media Access Control (MAC) address of VM adapters and Host adapters.
- **P-Mode** — provides a check box for all adapters on either the local or remote computer where you can specify message capture in Promiscuous Mode with the **Microsoft-Windows-NDIS-PacketCapture** provider, simply by selecting a **P-Mode** check box for a corresponding adapter. Note that you will need to determine which Host adapters on the target machine support Promiscuous Mode capture.

Finding Column Data
The tree grid in the Interface Selection section of the dialog also contains a **Show Column Filter Row** icon to the left of the **Name** column that enables you to search for text in the **Name** or **MAC Address** columns. This can be useful if you have a particularly long list of VM and host adapters. If you click the **Show Column Filter Row** icon, the **Column Filter Row** displays two amber-colored text boxes where you can enter search text. When you specify search text, the column is automatically filtered to display only those rows of data where a text match is found. You can remove the filtered view by clicking the **x** mark in the amber-colored text box above the column where you searched for text. This feature can help you to quickly isolate an adapter on which to capture data.

Configuring Host Adapter and Hyper-V-Switch Filters

The Filters section of the **Advanced Settings** dialog contains numerous settings in the **Filters** pane. The configuration settings that follow enable you to filter the data that you capture remotely on adapters that you select in the Interface Selection section of the dialog, although the impact of some filters can be different on host adapters versus the effects on VM adapters:

- **Layer** — this setting has a different meaning when applicable to capturing data on a Host adapter versus a Switch adapter, as follows:
    - **Host Adapter** — when you *select* the **All Layers** check box prior to running a remote trace, all packet traffic that is intended for host adapters will be intercepted on all NDIS layers, for traffic that traverses both up and down the filter stack. If you *unselect* the **All Layers** check box, then packets will be intercepted on the default/lowest NDIS layer only, in both traffic directions. See the article NDIS Filter Drivers on MSDN for an overview and illustration of the NDIS layers.

      In the diagram of the indicated article, the NDIS layers are shown as Filter Modules, which can also be thought of as nuances of the Link Layer. For example, some of the tasks that these Filter Modules can carry out include quality of service (QoS) processing that prioritizes packets and translation of Ethernet headers to Wi-Fi for wireless packets.

      For example, you might select the **All Layers** check box when you are investigating dropped packets, so that you can monitor all NDIS layers and determine which NDIS filter layer is dropping packets. In addition, because packets can be intercepted both down and up the NDIS filter stack, it is recommended for host adapters that you select both the **Ingress** and

Egress check boxes for the **Direction** parameter in **Advanced Settings**, which correspond to outbound packets (sent down the stack) and inbound packets (received up the stack), respectively. For more information about the **Ingress** and **Egress** options, see the "Direction" item in this list.

> **NOTE**
> If you happen to select all host adapters and VM adapters in the Interface Selection section of the **Advanced Settings** dialog and then run a remote trace, the functionality that is enabled by selecting the **All Layers**, **Ingress**, and **Egress** check boxes is automatically applied in the way that is appropriate for each of host adapters and VM adapters. However, for the sake of better interpretation of results, you might consider capturing on either host adapters or VM adapters, but not both at the same time.

- **Switch Adapter** — when you *select* the **All Layers** check box prior to collecting data on a remote VM adapter, all packet traffic passing through the Hyper-V-Switch that services the VM will be intercepted on all layers of the switch Extension stack in the Egress or/and Ingress path/s that you specify. If you *unselect* the **All Layers** check box, then by default, packets will only be intercepted on the highest filter layer of the Hyper-V-Switch extension stack in the Egress or/and Ingress path/s, as appropriate, during your remote trace. See the article Overview of the Hyper-V Extensible Switch on MSDN for a conceptual overview of the Hyper-V-Switch and its filtering stack.

  In the diagram provided in the indicated article, the filtering layers of the switch are shown as Capturing, Filtering, and Forwarding extensions, which can also be thought of as nuances of the Link Layer. The Hyper-V-Switch extension filter stack can inspect, monitor, drop, exclude, and duplicate packets; it also determines the packet source and destination switch ports and can enforce packet security and VM network policies. To help you better understand how to configure filtering in the **Advanced Settings** dialog when capturing traffic on a remote VM that is serviced by a remote Hyper-V-Switch, the following explanation describes the path of packets entering and leaving the Hyper-V-Switch. Note that this path is identical whether packets are being sent to or from a VM. See the indicated article to reference the components that are described here:

  A packet from a Child Partition VM or network adapter is sent to a Hyper-V-Switch port and then down the Ingress data path through the switch extension stack. Packets are processed through the extension stack filtering rules, while switch *source* port information is acquired in this path. The packet then goes up the Egress data path through the switch extension stack. Packets are processed through the extension stack filtering rules again, while switch *destination* port information is acquired in this path.

  > **NOTE**
  > Switch source port information is available on both the Ingress and Egress data paths, while switch destination port information is available on the Egress data path only.

  As a usage example, when you are simply monitoring traffic for a VM that is serviced by a Hyper-V-Switch, it is recommended that you select the **Egress** option only (and leave the **All Layers** option unselected) in the **Advanced Settings** dialog, so that packets will be intercepted on the default/highest extension stack layer in the Egress data path only. These settings can result in performance and usability improvements, as follows:

  - Bandwidth consumption is moderately decreased because packets are intercepted in

one path and on one extension only.

  - Filtering logic is more accurately executed because both source and destination port information is available in the Egress data path.
  - Message analysis is made simpler with lower packet count.

However, in troubleshooting scenarios where packets are being dropped and you need to investigate where this is occurring, you should select the **All Layers** option and enable both **Ingress** and **Egress** options to determine whether packets are being dropped in either of these paths in the Hyper-V-Switch extension stack. Note that the check boxes for both of these paths are enabled by default in the **Advanced Settings** dialog.

- **Truncation** — enables you to truncate packets that you capture from a remote host. You can truncate packets to limit bandwidth consumption and reduce the memory footprint of the traffic you capture, so that you receive only a certain amount of bytes for every packet. In practice, you might want to look at only the headers of a particular protocol of interest. If you know the length of the protocol header in bytes, you can set that value in the **Truncation** text box and receive only the header portion of the packets of interest in your trace. The default **Truncation** value that displays when you open the **Advanced Settings** dialog varies as indicated in the note below; however, you can set this value as required. Note that a value of zero (0) bytes means packets will not be truncated.

  > **IMPORTANT**
  >
  > For **Trace Scenarios** that use the **Microsoft-Windows-NDIS-PacketCapture** provider for *local* traces, the default **Truncation** value will be set to zero (**0**) bytes after you open the **Advanced Settings** dialog, while for **Trace Scenarios** that use the **Microsoft-Windows-NDIS-PacketCapture** provider for *remote* traces, the default **Truncation** value will set to **128** bytes after you open the **Advanced Settings** dialog. In either of these cases, if you *set* the **Truncation** value to a number greater than zero (**0**) bytes, Message Analyzer will only parse a limited subset of protocols, which includes the Ethernet, GRE, IPv4, IPv6, ESP, AH, IKE, AIPS, TCP, UDP, and HTTP protocols.

- **Direction** — when capturing traffic on host adapters, the **Ingress** and **Egress** options enable you to specify the direction in which to capture packets that traverse the NDIS filter stack — either outbound traffic that corresponds to the **Ingress** option and down the stack, or inbound traffic that corresponds to the **Egress** option and up the stack. When capturing traffic on a VM that is serviced by a Hyper-V-Switch, these same options specify the path that packets take when traversing the Hyper-V-Switch extension stack, for *both* inbound and outbound traffic to and from a VM. For recommendations on how to set these options, see the "Host Adapter" and "Switch Adapter" list items above.

- **EtherTypes** — enables you to specify an EtherType number that represents a protocol, such as IPv4, IPv6, or ARP, that is in the payload of an Ethernet frame. For example, if you specify an EtherType hexadecimal value in the **EtherType** text box for one or more of these protocols, comma-delimited and without the “0x” designator—for example `0800`, `86DD`, `0806`, respectively — the remote trace will filter for and return only Ethernet frames that have IPv4, IPv6, or ARP in their payloads.

> **NOTE**
>
> It is recommended that you always specify an **EtherType** when you are configuring an **IP Protocol Number** or **IP Address** filter. By specifying an **EtherType** value in these cases, you ensure that only the targeted traffic is returned in the remote trace.
>
> For a list of EtherTypes, see IEEE 802 Numbers.

- **IP Protocol Numbers** — enables you to specify an **IP Protocol Number** that filters for IP packets that have their **Protocol** header field set to the protocol number that you specify. Moreover, it filters for IP packets with a protocol in the payload that is identified by the specified protocol number, which is commonly TCP, UDP, and ICMP. For example, if you specify an IP protocol number in hexadecimal format in the **IP Protocol Numbers** text box for one or more of these protocols, comma-delimited and without the "0x" designator — for example `06`, `11`, `01`, respectively — the remote trace will filter for and return only IP packets that have TCP, UDP, or ICMP in their payloads.

  > **NOTE**
  >
  > It is strongly recommended to always specify an EtherType when configuring IP protocol number filters. For example, if you specify a TCP protocol number without also providing an EtherType number for IPv4 or IPv6, the filter will be applied to IP packets only. However, all other non-IP based traffic will be passed as well, which can be an unexpected result. To avoid this result in the example, you can specify the Protocol Number (`06`) for TCP and the EtherType number (`0800`) for IPv4.

  For a list of Protocol Numbers, see Protocol Numbers.

- **Mac Addresses** — enables you to create a filter that specifies one or more MAC addresses. The format consists of dash-delimited, Link Layer addresses such as `10-60-4B-6D-8D-2D`. You can use **Mac Address** filters to control the adapter on which you capture traffic. For convenience, the MAC addresses of all adapters enumerated on the remote host are provided in the tree grid section of the **Advanced Settings** dialog. Even if you have all adapters selected in **Advanced Settings**, you can use **Mac Address** filtering to limit the remote traffic you capture to the adapter/s with the addresses you specify.

  > **IMPORTANT**
  >
  > If you want to capture traffic from a specific remote VM, you will need to select the VM in the tree grid section of the **Advanced Settings** dialog and then create a **MAC Address** filter based on the VM's MAC address to isolate the data. Otherwise, you will capture Hyper-V-Switch traffic that is destined for all VMs that are serviced by the switch, given that a Hyper-V-Switch driver cannot of itself distinguish between VMs.

- **IP Addresses** — enables you to create a filter that specifies one or more IPv4 or IPv6 addresses. You enter IPv4 and IPv6 addresses in the standard format for these types of addresses and comma-delimit them, for example, `192.168.1.1`, `192.168.1.1`, and so on. You can use an **IP Address** filter to isolate and return remote traffic from a particular machine that is assigned the IP address that you specify. However, if you want to return *only* packet traffic for the specified IP address, it is recommended that you also specify an EtherType for the IP address that you provide; for example, specify an IPv4 EtherType when filtering on an IPv4 address. If you specify an IPv4 address alone without also providing an EtherType number, the filter is applied to IPv4 packets only, but all other non-IP–based traffic will be unfiltered and therefore passed, which can be an unexpected result. To avoid this result in the example, you can specify the IPv4 address in the

format `192.168.1.1` and include the EtherType number `0800` for IPv4.

See Also

Configuring a Remote Capture

Performing a Live Capture

4 minutes to read

This topic describes several methods that you can use to perform a live capture with Message Analyzer. It also describes how to stop, pause, resume, or restart a Live Trace Session.

Methods for Starting a Live Trace Session

The first two methods in the list that follows require no initial configuration, which allows you to start a Live Trace Session with as little as a single click. The third method allows you to create a custom capture configuration before you start a Live Trace Session, which is a little more involved.

- **Start Local Trace** — click the **Start Local Trace** button on the Message Analyzer **Start Page**. Use this method to instantly start a trace that runs on your local computer and captures packets from the network with the **Microsoft-PEF-NDIS-PacketCapture** provider or the **Microsoft-Windows-NDIS-PacketCapture** provider, depending on the operating system that is running on the local computer, as described by the table in Built-In Trace Scenarios.
- **Favorite Scenarios** — click a **Favorite Scenario** on the Message Analyzer **Start Page**. Use this method to select one of the following default **Favorite Scenarios** and instantly start capturing data at the Data Link Layer, Transport Layer, or Application Layer, respectively:
  - **Local Network Interfaces** — uses the **Microsoft-PEF-NDIS-PacketCapture** provider on computers running the Windows 7, Windows 8, or Windows Server 2012 operating system. Uses the **Microsoft-Windows-NDIS-PacketCapture** provider on computers running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system.
  - **Loopback and Unencrypted IPSEC** — uses the **Microsoft-PEF-WFP-MessageProvider** for any supported operating system that is running on the local computer.
  - **Pre-Encryption for HTTPS** — uses the **Microsoft-Pef-WebProxy** and Fiddler providers for any supported operating system that is running on the local computer.

You will find each of the above **Trace Scenario** Favorites by performing any of the following actions:

```
-   Click the Message Analyzer **File** menu, point to the **Favorite Scenarios** item, and then select one
of the above scenarios in the submenu that appears.

-   Click the **Favorite Scenarios** drop-down list on the Message Analyzer global toolbar, and then select
one of the above scenarios.

-   Click one of the scenarios in the **Favorite Scenarios** list on the **Start Page**.

 > [!TIP]
 >  To add more scenarios to the **Favorite Scenarios** list, which includes other built-in **Trace
Scenarios** and any custom scenarios that you have created, click **Edit Favorites** on the **Start Page**
to display the **Edit Favorites** dialog. From this dialog, configure a **Trace Scenario** as a Favorite and
add it to the **Favorite Scenarios** list by clicking the white star next to a **Trace Scenario**, at which
time the white star changes to the color yellow and the **Favorite Scenarios** list is updated to include
the new Favorite that you are adding. The update is then reflected in the previously indicated locations
where you can access the **Favorite Scenarios** list.
```

- **New Session** dialog — click the **Start** button in the **New Session** dialog. Use this method to create the capture configuration for a Live Trace Session prior to starting it. You could configure a session simply by

selecting a particular **Trace Scenario** and then starting the session. Optionally, you can customize the capture configuration first with selected system ETW providers; with filtering configurations that can include a **Fast Filter**, **Session Filter**, **Parsing Level**, and an event **Keyword** bitmask or error **Level** filter; with advanced host adapter and/or switch adapter filters; by specifying a data viewer; and so on. If you want to create unique capture configurations that are tailored to your environment or other requirements, this is the method you will use most often.

> **NOTE**
>
> After you create a custom capture configuration for a Live Trace Session, you can save the configuration as a new user **Trace Scenario**. It will then appear in the **My Items** category of the **Select Scenario** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog, as described in Saving Trace Scenarios. You can then run the scenario on demand. Note that you can fully edit or delete any **Trace Scenarios** that you created, whereas, this is not possible for any of the built-in **Trace Scenarios** that install with Message Analyzer.

More Information

To learn more about how to create a unique capture configuration for a Live Trace Session, see Configuring a Live Trace Session.

Managing Session Stop, Pause, Resume, and Restart

While a Live Trace Session is in progress, you have the option to Pause the session by clicking the **Pause/Resume** button on the global Message Analyzer toolbar. You might do this at a point in time where you need to manually trigger some process, application, or event for which you want to capture messages. Immediately after you initiate such a trigger, you can resume the Live Trace Session by clicking the **Pause/Resume** button again.

When your data capture is complete, you can Stop the Live Trace Session by clicking the **Stop** button on the global Message Analyzer toolbar. In addition, for any Live Trace Session that is Stopped, you can Restart the session by clicking the **Restart** button on the global Message Analyzer toolbar. Note that you cannot Restart a session that is in the Paused state, as you can only Resume or Stop a session that is Paused. Also be aware that when you Restart a Live Trace Session, all existing captured data will be lost unless you save it before restarting.

Procedures: Using the Network Tracing Features

29 minutes to read

The procedures in this section encapsulate some of the main functionalities described in the Capturing Message Data section, which includes defining the scope of data capture in a Live Trace Session. Although you can quickly start a Live Trace Session with a single click of the **Start Local Trace** button or of a specific **Favorite Scenario** on the Message Analyzer **Start Page**, you might want to specify your own Live Trace Session configuration settings before starting a trace. You can do this by clicking the **New Session** button on the **Start Page** to open the **New Session** dialog, from where you can specify a **Live Trace** as a data source and then customize the capture configuration of your Live Trace Session. You can also access the same configuration settings for a Live Trace Session by clicking the **New Session** button in the upper left corner of the Message Analyzer UI, or by selecting the **New Session** item from the Message Analyzer **File** menu.

You will use the **New Session** dialog to create capture configurations in all of the procedures in this section. Each of the procedures specified in this section are preceded by conceptual information that describes the purpose, configuration features, and/or expected results of the procedure.

> **NOTE**
> The procedures serve as simple examples that are intended to demonstrate how to utilize Message Analyzer tracing configurations and other facilities, rather than providing a comprehensive treatment of network troubleshooting. Although these procedures demonstrate the use of Message Analyzer capabilities in some basic scenarios, they are only a sampling of what you can accomplish with Message Analyzer, given that you can also apply the methodologies described here to many other scenarios.

Procedure Overviews

A brief description of each procedure is included here for review, as follows.

Configure and Run a Local Network Interfaces Trace — provides an example of how to modify the default **Local Network Interfaces Trace Scenario**; by specifying an adapter on which to capture messages and by adding a combination of filters to the **Microsoft-PEF-NDIS-PacketCapture** provider configuration on computers running the Windows 7, Windows 8, or Windows Server 2012 operating system; or to the **Microsoft-Windows-NDIS-PacketCapture** provider configuration on computers running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system; that restrict the scope of data retrieval to only messages that pass the defined filtering criteria.

Configure and Run a Loopback and Unencrypted IPSEC Trace — provides an example of how to modify the default **Loopback and Unencrypted IPSEC Trace Scenario** by setting the **Microsoft-PEF-WFP-MessageProvider** configuration to capture only HTTP packets through a TCP port filter.

Configure and Run a Pre-Encryption for HTTPS trace — provides an example of how to modify the default **Pre-Encryption for HTTPS Trace Scenario** by defining filtering criteria that enables you to monitor HTTP message exchanges between a client browser and a specified web server.

Capture Traffic on a Remote Host — provides an example of how to use the default **Remote Network Interfaces Trace Scenario** to capture data on a remote host that is running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system. Includes specifying special filtering settings for a Hyper-V Switch and a target virtual machine (VM) that it services.

Design and Run a Custom Trace Scenario — provides an example of how to create, save, and run a **Trace**

Scenario template that monitors the manual Group Policy update process on the local machine for signs of any issues with Lightweight Directory Access Protocol (LDAP) communications.

> **IMPORTANT**
> If you have not logged off Windows after the first installation of Message Analyzer, please log off and then log back on before performing these procedures. This action ensures that in all subsequent logons following installation, your security token will be updated with the required security credentials from the Message Capture Users Group (MCUG). Otherwise, you will be unable to capture network traffic in local **Trace Scenarios** that use the **Microsoft-Windows-NDIS-PacketCapture** provider, unless you start Message Analyzer with the right-click **Run as administrator** option.

> **NOTE**
> Even if you log off your system, log back on, and receive the required security credentials from the MCUG, you will still need to use the **Run as administrator** option to capture message data with the **Microsoft-Windows-NDIS-PacketCapture** provider in the procedure Capture Traffic on a Remote Host. This is the result of the inherent remote capabilities of this provider and the security restrictions that must therefore be applied to it.

Configure and Run a Local Network Interfaces Trace

In the following procedure, you will select the default **Local Network Interfaces Trace Scenario** on a computer that is running the Windows 7, Windows 8, or Windows Server 2012 operating system. You will then configure the **Microsoft-PEF-NDIS-PacketCapture** provider to isolate captured messages to a particular network adapter device and a specific IPv4 address. You might use a trace configuration such as this to minimize disk and CPU impact while capturing data on a busy computer that is overwhelmed with traffic.

> **NOTE**
> If you are running the Window 8.1, Windows Server 2012 R2, or Windows 10 operating system, and you select the **Local Network Interfaces Trace Scenario** in the procedure that follows, the **Microsoft-Windows-NDIS-PacketCapture** provider used in this scenario has different filtering options than the **Microsoft-PEF-NDIS-PacketCapture** provider. For more information about how to specify filters for the **Microsoft-Windows-NDIS-PacketCapture** provider, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog. In addition, you will need to run Message Analyzer with the right-click **Run as administrator** option on computers running one of the above specified operating systems, due to security restrictions of the **Microsoft-Windows-NDIS-PacketCapture** provider.

To configure and run a Local Network Interfaces trace

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.
2. On the Message Analyzer **Start Page**, click the **New Session** button to display the **New Session** dialog.
3. Under **Add Data Source** in the **New Session** dialog, click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features that it contains in the **New Session** dialog.
4. In the **Network** category of the **Select Scenario** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog, click the **Local Network Interfaces Trace Scenario**.

   If your operating system is Windows 7, Windows 8, or Windows Server 2012, the **ETW Providers** list on the **Live Trace** tab is populated with the **Microsoft-PEF-NDIS-PacketCapture** provider **Name** and **Id** (GUID). Otherwise, for the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system, the **Microsoft-Windows-NDIS-PacketCapture** provider information displays.
5. In the **ETW Providers** list on the **Live Trace** tab, ensure that the **Microsoft-PEF-NDIS-PacketCapture**

provider is selected and then click the **Configure** link to the right of its **Id** to display the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, as shown in Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog. If you are working with the **Microsoft-Windows-NDIS-PacketCapture** provider, clicking the **Configure** link will display the **Advanced Settings – Microsoft-Windows-NDIS-PacketCapture** dialog, as shown in Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

6. If you are working with one of the specified earlier operating systems and the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, proceed to the next step. Otherwise, if you are working with one of the specified later operating systems and the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog, proceed to step 15.

7. In the **System Network** tree grid on the **Provider** tab of the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, specify a physical or wireless adapter on which to capture data, by selecting the **In** and **Out** direction check boxes of the adapter.

   The **Microsoft-PEF-NDIS-PacketCapture** provider is set to capture both inbound and outbound traffic on the adapter device that you specified.

8. In the **Fast Filters** pane of the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog under **Group 1**, click the drop-down arrow next to the **Fast Filter 1** designator to display the filter type menu items and then select the **IPv4Address** filter type from the menu.

9. In the text box to the right of the filter type drop-down, enter an IPv4 address in a format similar to the following:

   192.168.1.1

10. In the **Fast Filters** pane of the **Advanced Settings** dialog under **Group 2**, click the drop-down arrow next to the **Fast Filter 1** designator to display the filter type menu items and then select the **LinkLevelAddress** filter type from the menu.

11. In the text box to the right of the filter type drop-down, enter a MAC address for a different adapter that you want to block traffic to, in a format similar to the following:

    !=00-2F-39-7E-1F-36

    This filter ensures that traffic will be blocked from reaching the adapter for which you specified the negated **LinkLevelAddress**. Note that you can also achieve this same result by simply deselecting the **In** and **Out** directional check boxes on the adapter for which you want to block traffic. However, this example shows you a simple way to utilize filter **Groups**.

12. In the **Advanced Settings** dialog, highlight the **System Network** tree grid row that contains the adapter device you initially specified and then click the **Apply To Highlighted** button in **Group 1** of the **Fast Filters** pane to assign the filter **Group** to the adapter.

> **NOTE**
> When you click the **Apply To Highlighted** button, the name of the adapter device to which the **Fast Filter Group** is applied appears next to the **Target** label for the corresponding **Group**.

13. In the **Advanced Settings** dialog, highlight the **System Network** tree grid row that contains the adapter device for which you specified a negated **LinkLevelAddress** filter and then click the **Apply To Highlighted** button in **Group 2** of the **Fast Filters** pane to assign the filter **Group** to the adapter.

    The **Microsoft-PEF-NDIS-PacketCapture** provider is now configured to do the following in your Live Trace Session:

- Isolate trace data to only the adapter device that you initially specified.
- Block all packets to the device for which you created a negated **LinkLevelAddress** filter.
- Target data for a specific IPv4 address.
- Reduce message count and improve trace performance.

When packets arrive that are intended for the adapter device that you initially specified, the filter configuration for **Group 1** is applied to those packets to pass the message data. When packets arrive that are intended for the second adapter device, the filter configuration for **Group 2** is applied to those packets to block the message data.

14. Click **OK** to exit the **Advanced Settings – Microsoft-PEF-NDIS-PacketCapture** dialog.
15. If you are working with one of the specified later operating systems and the **Microsoft-Windows-NDIS-PacketCapture** provider, select the check box in the tree grid (Interface Selection) section of the **Advanced Settings – Microsoft-Windows-NDIS-PacketCapture** dialog for the Ethernet or wireless adapter on which to capture data, and unselect the other check boxes. Otherwise, if you are working with the **Microsoft-PEF-NDIS-PacketCapture** provider, proceed to step 19.
16. In the **EtherTypes** text box of the **Advanced Settings** dialog, enter the Ethernet type value for an IPv4 address, as follows:
    0800
17. In the **IP Addresses** text box of the **Advanced Settings** dialog, enter the value of the IP address of the local computer in a format similar to the following:
    192.168.1.1

    The **Microsoft-Windows-NDIS-PacketCapture** provider is now configured to do the following in your Live Trace Session:

    - Isolate packet traffic to only the adapter device for which you selected a check box.
    - Block traffic to all other adapter devices.
    - Capture packet traffic for the target IPv4 address only, while removing all other traffic with the specified EtherType value.
    - Reduce message count and improve trace performance.
18. Click **OK** to exit the **Advanced Settings – Microsoft-Windows-NDIS-PacketCapture** dialog.
19. In the **New Session** dialog, you can optionally enter a name for the session in the **Name** text box.
20. If the **Analysis Grid** is not already specified as the data viewer for your Live Trace Session, click the **Start With** drop-down menu in the **New Session** dialog and select it.
21. Click the **Start** button in the **New Session** dialog to begin capturing data in your Live Trace Session.

    Message Analyzer may begin capturing data immediately.
22. As captured messages accumulate in a new **Analysis Grid** session viewer tab below the Message Analyzer global toolbar, attempt to reproduce any conditions that are related to a particular issue you might be having on the target computer.
23. Stop the trace at a suitable point by clicking the **Stop** button on the global Message Analyzer toolbar, or use the keyboard shortcut `Shift+F5`.
24. In the **Analysis Grid**, right-click the **DiagnosisTypes** column header and select **Group** from the menu that displays, to group any error messages you might have received, for further analysis of a related issue.

All Diagnosis messages of the same type will display in a separate group that is labeled by the **DiagnosisType**. If you have multiple types of Diagnosis messages, there will be a unique group for each type. You can expand these groups to view the messages associated with each type. For further details about Diagnosis message types, see the Diagnosis Category topic.

Configure and Run a Loopback and Unencrypted IPSEC Trace

In the following procedure, you will select the built-in **Loopback and Unencrypted IPSEC Trace Scenario** and configure a **Fast Filter** to retrieve data from **TCP port** 80, thereby filtering for HTTP traffic only. You might use a **Trace Scenario** such as this on a client computer to limit your capture to HTTP traffic only, along with the protocol stack that supports the HTTP operations. This can help you to troubleshoot webpage performance, detect issues with HTTP connectivity, or debug a website based on HTTP responses sent to the client. Also, the filter employed in this scenario minimizes the impact on disk I/O and the CPU because the filter selects specific messages for capture, resulting in reduced message count and thus better performance.

> **NOTE**
> In this scenario, the **TCP port** filter will pass messages that transit both TCP source and destination ports.

To configure and run a Loopback and Unencrypted IPSEC trace

1. Start Message Analyzer as indicated in the first procedure of this section.
2. On the Message Analyzer **Start Page**, click the **New Session** button to display the **New Session** dialog.
3. Under **Add Data Source** in the **New Session** dialog, click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features that it contains.
4. In the **Network** category of the **Select Scenario** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog, click the **Loopback and Unencrypted IPSEC Trace Scenario**.
5. In the **ETW Providers** list on the **Live Trace** tab, click the **Configure** link to the right of the **Id** for the **Microsoft-PEF-WFP-MessageProvider** to display the **Advanced Settings - Microsoft-PEF-WFP-MessageProvider** dialog, as shown in Using the Advanced Settings- Microsoft-PEF-WFP-MessageProvider Dialog.
6. In the **Fast Filters** pane on the **Provider** tab of the **Advanced Settings** dialog, click the **Fast Filter 1** drop-down arrow and select the **TCP port** item in the drop-down list.
7. In the text box to the right of the drop-down selection you made, enter the number *80*.

   The **Microsoft-PEF-WFP-MessageProvider** is now configured to filter for HTTP packets on TCP port 80.

   The messages that this trace configuration returns can include WFPCapture events, HTTP operations (as indicated by blue-cubed icons to the left of message numbers) the underlying message stack that supported such operations such as TCP packets captured on port 80, in addition to TCP fragments and other HTTP messages.
8. Click **OK** to exit the **Advanced Settings - Microsoft-PEF-WFP-MessageProvider** dialog.
9. In the **New Session** dialog, optionally specify a name for the Live Trace Session in the **Name** text box.
10. If the **Analysis Grid** is not already specified as the data viewer for your Live Trace Session, click the **Start With** drop-down list in the **New Session** dialog and select it.
11. Click the **Start** button in the **New Session** dialog to begin capturing data in your Live Trace Session.

    Message Analyzer may begin capturing data immediately.

12. As captured messages accumulate in a new **Analysis Grid** session viewer tab below the global Message Analyzer toolbar, attempt to reproduce any conditions that may be related to HTTP connectivity or performance problems, for example, by navigating to a web server where clients experience these issues.
13. Stop the trace at a suitable point by clicking the **Stop** button on the Message Analyzer global toolbar.
14. In the **Analysis Grid** viewer, right-click the column with the **DiagnosisTypes** icon and select **Group** from the context menu that displays to group any diagnostic messages you might have received, for further analysis.
15. Review HTTP **StatusCodes** for evidence of connection or performance issues on the server, as described in Addendum 2: HTTP Status Codes of this documentation.

> **NOTE**
> To create a prominent view of HTTP status data, add the **HTTP.Response.StatusCode** field to the default **Analysis Grid** viewer column **Layout** with the **Field Chooser Tool Window**, by right-clicking the **StatusCode** field name and selecting the **Add As Column** command from the context menu that appears, as described in Using the Field Chooser.

Configure and Run a Pre-Encryption for HTTPS trace

In the following procedure, you will run the **Pre-Encryption for HTTPS Trace Scenario** on a client computer with a filter configuration that enables you to capture and monitor unencrypted HTTP browser messages that are sent to a specified HTTP host that is slow to send response messages.

To configure and run a Pre-Encryption for HTTPS trace

1. Start Message Analyzer as indicated in the first procedure of this section.
2. On the Message Analyzer **Start Page**, click the **New Session** button to display the **New Session** dialog.
3. Under **Add Data Source** in the **New Session** dialog, click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features that it contains.
4. In the **Network** category of the **Select Scenario** drop-down list on the **ETW Providers** toolbar, click the **Pre-Encryption for HTTPS Trace Scenario**.
5. In the **ETW Providers** list on the **Live Trace** tab, click the **Configure** link to the right of the **Microsoft-Pef-WebProxy** provider **Id** to display the **Advanced Settings - Microsoft-Pef-WebProxy** dialog.
6. On the **Provider** tab of the **Advanced Settings** dialog, specify the host name for a web server that may be slowly responding in the text box to the right of the **Hostname Filter** property of the provider, in a format similar to the following:

   www.xxxxx.com.
7. On the **Provider** tab of the **Advanced Settings** dialog, specify an HTTP port number in the text box to the right of the **Port Filter** property of the provider, to ensure that you capture HTTP traffic only. Specify the port number in integer format, as indicated in the following examples:

   80 for HTTP, or *443* for HTTPS

   The **Microsoft-PEF-WebProxy** provider is now configured to capture HTTP packets that are sent to and from the specified web server.
8. Click **OK** to exit the **Advanced Settings - Microsoft-PEF-WebProxy** dialog.
9. In the **New Session** dialog, optionally specify a name for the Live Trace Session in the **Name** text box.
10. If the **Analysis Grid** viewer is not already specified as the data viewer for your Live Trace Session, click the

Start With drop-down list in the **New Session** dialog and select it.

11. Click the **Start** button in the **New Session** dialog to begin capturing data in your Live Trace Session.

    Message Analyzer may begin capturing data immediately.

12. As captured messages accumulate in a new **Analysis Grid** session viewer tab below the global Message Analyzer toolbar, open a web browser and establish a connection to the specified HTTP host by launching an HTTP request to the associated site address.

13. Stop the trace at a suitable point by clicking the **Stop** button on the Message Analyzer global toolbar.

14. In the **Analysis Grid** viewer, right-click the column with the **DiagnosisType** icon and select **Group** from the context menu that displays, to group any diagnostic messages you might have received, for further analysis.

15. Review HTTP **StatusCodes** for evidence of connection or performance issues on the server, as described in Addendum 2: HTTP Status Codes of this documentation.

    To create a prominent view of HTTP status data, add the **HTTP.Response.StatusCode** field to the default **Analysis Grid** viewer column **Layout** with the **Field Chooser** dialog, by right-clicking the **StatusCode** field name and selecting the **Add As Column** command from the context menu that appears, as described in Using the Field Chooser.

> **TIP**
> You can also **Group** the **StatusCode** column in the **Analysis Grid** viewer to organize status codes into groups, for ease of analysis. To do so, right-click on the **StatusCode** column and select the **Group** command from the context menu that appears.

Capture Traffic on a Remote Host

In the following procedure, you will use the **Remote Network Interfaces Trace Scenario** on a computer that is running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system, to capture traffic from a virtual machine (VM) that is serviced by a Hyper-V-Switch on a remote computer that is running one of the same operating systems. In the procedure, you will do the following:

- Select the **Remote Network Interfaces Trace Scenario**.
- Specify the remote host connection credentials.
- Make an initial connection with the host to enumerate its adapter configuration.
- Use the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog to specify special filtering configurations for the Hyper-V-Switch and the VM from which you will capture remote message traffic.

To configure and run a Remote Network Interfaces trace

1. Start Message Analyzer as indicated in the first procedure of this section.

2. On the Message Analyzer **Start Page**, click the **New Session** button to display the **New Session** dialog.

3. Under **Add Data Source** in the **New Session** dialog, click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features that it contains.

4. In the **Network** category of the **Select Scenario** drop-down list on the **ETW Providers** toolbar, click the **Remote Network Interfaces Trace Scenario**.

5. On the **Live Trace** tab of the **New Session** dialog, click the **Edit** button adjacent to the **Target Computers**

list to display the **Edit Target Computers** dialog.

6. Click the **Add** drop-down arrow in the **Edit Target Computers** dialog and select the **New Row** item in the menu that displays.

   A new row is added to the target computers grid in the dialog.

7. Specify the name or IP address of the remote host on which you intend to capture message traffic, by entering it in the new row under the **Computer Name/IP Address** column of the **Edit Target Computers** dialog. You can specify the remote host name by simply entering the host name without including the forward backslashes that are customarily used in the Universal Naming Convention (UNC) standard.

> **TIP**
> You have the option to capture traffic on multiple remote hosts, as long as each one is running a supported operating system. For each remote host, you will need to create a new row in the **Edit Target Computers** dialog and specify the host name and connection credentials. When your trace results are complete, all the captured message data will be returned to Message Analyzer and aggregated into a single **Analysis Grid** session viewer tab.

8. If you cannot use your current logon credentials to connect with the remote host, then specify an appropriate **User Name** and **Password** in the indicated columns of the new row you added. When specifying other logon credentials, use the Domain\Username format. Otherwise, you can leave these grid fields blank to connect to the remote host with your current logon credentials.

9. In the grid of the **Edit Target Computers** dialog, optionally select the row that contains the default **Localhost** setting and then click **Delete** on the dialog toolbar.

   The target computer configuration is now set to capture message traffic on the specified remote host only, assuming that you opted to remove the **Localhost** from the **Edit Target Computers** dialog.

10. When complete, click **OK** to exit the **Edit Target Computers** dialog.

11. In the **ETW Providers** list on the **Live Trace** tab, click the **Configure** link to the right of the **Microsoft-Windows-NDIS-PacketCapture** provider **Id** to display the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog.

12. On the **Provider** tab of the **Advanced Settings** dialog, click the **Host** drop-down arrow and select the name of the remote host in the drop-down list. This drop-down appears only if there is more than one host from which you are capturing traffic, for example, a remote host and the local computer, or one or more remote hosts.

    Message Analyzer attempts to connect with the remote host to enumerate the host and/or switch adapters on that computer. When complete, the enumerated adapters, switches, and VMs that Message Analyzer discovered on the remote host will populate the tree grid section of the **Advanced Settings** dialog.

13. In the tree grid section of the **Advanced Settings** dialog, remove all selected adapters, switches, and VMs from configuration by deselecting the **Machine** check box.

14. In the tree grid section of the **Advanced Settings** dialog, specify the VM on which to capture data by selecting the enabling check box in the second grid column for the target VM.

15. For the **Layer** parameter in the **Filters** pane of the **Advanced Settings** dialog, select the **All Layers** check box to ensure that packets are intercepted at all Hyper-V-Switch extension layers and so that the filtering rules of each switch extension are applied to all packets that traverse the switch layers.

16. For the **Direction** parameter in the **Filters** pane of the **Advanced Settings** dialog, select the **Egress** check box so that packets are intercepted on all Hyper-V-Switch extension layers, but only in the direction that

you specified, which in this case is the egress path that goes up the switch extensions stack. Note that the ingress path goes down the extension stack.

Selecting **Egress** only will result in faster switch port management and subsequently an improvement in performance.

17. For the **EtherType** parameter in the **Filters** pane of the **Advanced Settings** dialog, specify the hexadecimal value `0800`, without the "0x" designator, to target the IPv4 protocol.

18. For the **IP Protocol Numbers** parameter in the **Filters** pane of the **Advanced Settings** dialog, specify the hexadecimal value `06`, without the "0x" designator, to target the TCP protocol.

    The **EtherType** and **IP Protocol Number** settings that you specified will cause the remote trace to filter for and return only Ethernet frames that have IPv4 packet payloads, and of those IPv4 packets, only the ones that have TCP payloads.

19. For the **MAC Addresses** parameter in the **Filters** pane of the **Advanced Settings** dialog, specify the MAC address of the target VM in a format similar to the following, to ensure that your **Remote Network Interfaces** trace returns remote traffic for the target VM only:

    ```
    10-60-4B-6D-8D-2D
    ```

20. Click **OK** to exit the **Advanced Settings** dialog.

21. Start your remote Live Trace Session by clicking the **Start** button in the **New Session** dialog.

    Message Analyzer may begin capturing data immediately.

22. As remote traffic from the specified VM begins to accumulate in the **Analysis Grid** viewer, perform operations on the remote VM or attempt to reproduce any issues that may be occurring on the target VM, or on the Hyper-V-Switch that services it.

    For example, you may be concerned with packets being lost from a particular protocol on the VM. Because you have enabled all extension layers of the Hyper-V-Switch to intercept packets, then if any packets are being dropped by a switch extension layer, they should generate events that you can detect in Message Analyzer trace results.

23. Stop the remote trace at a suitable point by clicking the **Stop** button on the Message Analyzer global toolbar, so that you can analyze your data.

24. Search for dropped packets, if you suspect that this is occurring in a Hyper-V-Switch extension layer. Do this by looking for ETW messages that contain the **ut:Dropped** event by applying the following view **Filter** from the Filtering toolbar that is located just above the **Analysis Grid** viewer:

    ```
    Etw.EtwProviderMsg.EventRecord.Header.Descriptor.Keywords == 0x0000010000000000
    ```

    You can also check to see if the KW_DROPPED flag is set in the **Flags** field of any ETW message in the **Details Tool Window**.

> **NOTE**
>
> To make it easier to analyze ETW messages, select the **ETW Layer Viewpoint** from the **Viewpoints** drop-down list on the Filtering toolbar to display all ETW messages with no layers above them.

TIP

When analyzing data that you have captured from multiple remote computers, you have the option to organize and summarize the captured data into groups that are labeled by host (data source) name. You can do this by adding the **DataSource** field from the **General** category of the **Field Chooser** to the default **Analysis Grid** viewer column **Layout**, and then applying the **Group** command by selecting it from the context menu that displays after you right-click the newly added **DataSource** column.

More Information

To learn more about the extension filtering stack on a Hyper-V-Switch, see Overview of the Hyper-V Extensible Switch on MSDN.

To learn more about capturing traffic on a remote host and specifying adapter and filter configurations for the **Microsoft-Windows-NDIS-PacketCapture** provider, see Configuring a Remote Capture.

To learn more about the **Field Chooser**, see Using the Field Chooser.

Design and Run a Custom Trace Scenario

In the following procedure, you will create a custom **Trace Scenario** template that captures LDAP traffic on the local client computer during a manual Group Policy update. You can run the template file whenever it is necessary to ascertain whether a client computer is experiencing Group Policy update issues.

This procedure primarily uses the **Loopback and Unencrypted IPSEC Trace Scenario** to take advantage of the capability of the **Microsoft-PEF-WFP-MessageProvider** to focus on messages above the Network Layer. However, in the procedure, you also have the option to add the **Microsoft-Windows-LDAP-Client** system ETW Provider to the trace configuration by specifying the **SASL LDAP Pre-encryption with WFP** scenario instead of the **Loopback and Unencrypted IPSEC** scenario, so that you can capture LDAP traffic unencrypted. In addition, the events written by the **Microsoft-Windows-LDAP-Client** provider can help you to better understand the state of the LDAP client when LDAP bind, search, request, and response messages are sent during Group Policy update.

To design and run a custom Trace Scenario

1. Start Message Analyzer as indicated in the first procedure of this section.
2. On the Message Analyzer **Start Page**, click the **New Session** button to display the **New Session** dialog.
3. Under **Add Data Source** in the **New Session** dialog, click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features that it contains.
4. In the **Network** category of the **Select Scenario** drop-down list on the **ETW Providers** toolbar, click the **Loopback and Unencrypted IPSEC Trace Scenario**.

TIP

If you want your custom scenario to capture LDAP traffic unencrypted, you might try using the **SASL LDAP Pre-encryption with WFP Trace Scenario**, which includes both the **Microsoft-Windows-LDAP-Client** and **Microsoft-PEF-WFP-MessageProvider** as part of the scenario configuration.

Note that you can optionally add other system ETW Providers to your **Trace Scenario** configuration from the **Add System Providers** dialog, which is accessible by clicking the **Add Providers** drop-down list on the **ETW Providers** toolbar. You might do this to return specific events that such a provider's **Keyword** configuration enables.

The **ETW Providers** list on the **Live Trace** tab is populated with the **Name** and **Id** (GUID) of the **Microsoft-PEF-WFP-MessageProvider** (and the **Microsoft-Windows-LDAP-Client** if you selected the **SASL LDAP Pre-encryption with WFP Trace Scenario**).

5. In the **ETW Providers** list, click the **Configure** link to the right of the **Id** for the **Microsoft-PEF-WFP-MessageProvider** to open the **Advanced Settings - Microsoft-PEF-WFP-MessageProvider** dialog, as shown in Using the Advanced Settings- Microsoft-PEF-WFP-MessageProvider Dialog.

6. In the **Fast Filters** pane on the **Provider** tab of the **Advanced Settings** dialog, click the drop-down arrow next to the **Fast Filter 1** designator to display the filter type menu items, and then select the **IPv4** filter type from the menu.

7. In the text box to the right of the filter type drop-down, enter an IPv4 address for the local computer in a format similar to the following:

   192.168.1.1

8. Click **OK** to exit the **Advanced Settings** dialog.

9. In the text box of the **Session Filter** pane in the **New Session** dialog, enter the following filter expression:

   ```
   *Port == IANA.Port.LDAP
   ```

   Your Live Trace Session template is now complete and configured to only capture LDAP traffic and other events related to the LDAP client, for the specified local IP address. In addition, the **Loopback and Unencrypted IPSEC Trace Scenario** and the **Session Filter** in use will remove a significant portion of lower-layer noise and improve performance.

> **TIP**
> To view the events that you can capture with the **Microsoft-Windows-LDAP-Client** ETW Provider, click the **Configure** link to the right of the **Id** for this provider to open the **Advanced Settings - Microsoft-Windows-LDAP-Client** dialog and then click the ellipsis (**...**) to the right of the **Keywords(Any)** or **Keywords(All)** text box. This action will display the **ETW Keyword Filter Property** dialog, from where you can view and select specific events to capture, that is, if they are triggered during a trace. For further information about setting **Keyword** bitmask filters, see System ETW Provider Event Keyword/Level Settings.

10. In the **New Session** dialog, optionally specify a name for your custom **Trace Scenario** in the **Name** text box.

11. On the **Live Trace** tab of the **New Session** dialog, click the **Save Scenario** button.

12. In the **Edit Trace Scenario** dialog that displays, provide a unique name for the scenario template in the **Name** text box and a description in the **Description** text box. Then choose an existing **Category** for the scenario template or specify a new one in the editable **Category** combo box.

13. Click the **Save** button in the **Edit Trace Scenario** dialog to save the scenario in the **Message Analyzer Trace Scenarios** Library and exit the dialog.

The **Trace Scenario** template that you saved should now display in the **My Items** category of the **Select Scenario** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog.

14. Display your **Trace Scenario** template configuration at any time by selecting it in the **Message Analyzer Trace Scenarios** Library that is accessible from the **Select Scenario** drop-down list.

When you do this, the **New Session** dialog will be populated with the custom settings that you specified when you created the **Trace Scenario** template. Note that you still have the option at this point to reconfigure your **Trace Scenario** prior to running a Live Trace Session; for example, you could specify a different **Session Filter**, provider line up, or **Keyword** bitmask configuration.

> **TIP**
> If you make further modifications to your **Trace Scenario** template, you can resave it with the new configuration settings without ever running it.

15. Start a Live Trace Session based on your custom **Trace Scenario** template by clicking the **Start** button in the **New Session** dialog.

    Message Analyzer may begin capturing data immediately.

16. While Message Analyzer is capturing message traffic, run the following command string from an elevated command prompt (Run as Administrator) to update Group Policy on the local machine:

    ```
    gpupdate /force
    ```

    The Live Trace Session begins capturing LDAP traffic on the local machine as the Group Policy update process accesses the appropriate Active Directory Group Policy Objects (GPOs) containing user and computer policy settings for the client.

17. Stop the trace at a suitable point by clicking the **Stop** button on the global Message Analyzer toolbar.

18. In the **Analysis Grid** viewer, right-click the **DiagnosisTypes** column and select **Group** from the menu that displays to group any diagnostic messages you might have received, for further analysis.

19. In the **Analysis Grid** viewer, review the LDAP messages for any status indications or errors that might reveal issues with LDAP bind, search, request, or response operations during Group Policy update. For example, you could add a **ResultCode** field as a new column to the default **Analysis Grid** viewer column **Layout** from the **Field Chooser Tool Window**.

Retrieving Message Data

2 minutes to read

This section provides conceptual details about the Message Analyzer Browse-Select-View (BSV) model and how you can employ its features for rapid and convenient management and processing of the target data that you will load into Message Analyzer through a Data Retrieval Session. It also describes the features that you can use to configure and run a Data Retrieval Session, along with various methods you can use to load saved data directly into Message Analyzer without additional configuration.

Go To Session Configuration

Go directly to an overview of Data Retrieval Session configuration workflow, filtering options, and other features that are available for configuring and starting a new Data Retrieval Session:

Configuring a Data Retrieval Session

What You Will Learn

In the topics of this section, you will learn how to accomplish the tasks indicated below.

Browse-Select-View Model — understand the BSV model and how to use it to configure a Data Retrieval Session.

Targeting Saved Data as an Input Source — read an introduction to starting, configuring, and optimizing a Data Retrieval Session that targets data from saved input files.

Configuring a Data Retrieval Session — examine a Data Retrieval Session configuration workflow example, filtering options, and the tasks you can perform when configuring a new Data Retrieval Session, such as:

- Locating supported input file types and targeting the saved data they contain as input to Message Analyzer.
- Decrypting data from saved traces.
- Using filtering to select the type of data to retrieve from saved files.
- Specifying a viewer in which to analyze the retrieved data.
- Opening text logs and creating OPN configuration files for parsing custom text logs.
- Handling cloud data as input to Message Analyzer from Azure storage tables or from Azure logs that are stored in binary large object (BLOB) containers.
- Aggregating and merging multiple related log and trace files for cross-file data correlation and analysis.
- Loading data from ETL files that contain WPP-generated events.

Performing Data Retrieval — use different methods to perform data retrieval.

Procedures: Using the Data Retrieval Features — perform example procedures that demonstrate the BSV model while encapsulating various aspects of data retrieval functionality.

More Information

To learn more about configuring session scenarios and reconfiguring an existing Data Retrieval Session, see the following topics:

Configuring Session Scenarios with Selected Data Sources

Editing Existing Sessions

Browse-Select-View Model

4 minutes to read

Microsoft Message Analyzer utilizes a unique set of Browse-Select-View (BSV) features that provide a versatile model for acquiring data, selecting specific data from the acquisition process, and presenting the data in various viewer formats. The BSV model allows you to navigate to saved log and trace files and to create a retrieval configuration that targets the data in specific files from which you choose to retrieve data. The model also enables you to extract subsets of message data from high volume trace files rather than extracting all data in target files, which can affect performance, impact memory usage, and make it difficult to work with the data. Because you can extract specific data from saved files, you save time upfront and clarify your troubleshooting efforts later on. The model also enables you to combine data from different sources, such as logs and trace files, and to view data from multiple sources in a single merged view.

The data selection aspect of the BSV model also applies to acquiring specific data through a live capture. For example, you can use various filters and different **Trace Scenarios** to isolate the precise data that you need to work with to quickly solve problems. In Message Analyzer, the goal of data selection is to acquire the least amount of data necessary to resolve an issue, in order to minimize consumption of system resources, improve performance, and streamline the data analysis process.

Acquiring Input Data

Although Message Analyzer makes use of two distinct paths in which to acquire input data for a session, Data Retrieval Sessions use the first of the following paths only:

- **Loading saved data** — this path handles historical/saved data that is loaded into Message Analyzer through a Data Retrieval Session.
- **Capturing live data** — this path handles new data that is captured in a Live Trace Session.

For each of these data acquisition paths, Message Analyzer uses common aspects of the BSV model to provide a similar data selection and viewing experience in both types of sessions. A Data Retrieval Session and a Live Trace Session also share another common characteristic, which is that they both enable you to acquire input from multiple data sources. For example, you can retrieve historical data from one or more saved source files in a single Data Retrieval Session; or you can capture data simultaneously from the local computer and/or multiple remote computers in a single Live Trace Session that establishes a subsession for each computer on which Message Analyzer is capturing data. Message Analyzer also has the capability to run multiple Live Trace Sessions at the same time, where you can start up multiple sessions one after the other.

The Message Analyzer BSV model supports the following operations when acquiring input data:

- **Browse for data sources** — you can browse for and load data from one or more data sources, such as saved trace and log files. See Targeting Saved Data as an Input Source for more information.

  The features of the BSV model not only facilitate loading data into Message Analyzer, but also enable you to select specific data that you want to load by using a **Session Filter** and/or a **Time Filter**, and also enable you to display data in several different viewer formats. You can also select an option that applies a pared-down parser set to input files that have truncated messages to improve performance and limit the types of messages retrieved. In addition, you can select configuration files that enable Message Analyzer to parse messages in text logs.

- **Capture message data** — you can capture new message data live by selecting one of the Message Analyzer default **Trace Scenarios**, as described in Built-In Trace Scenarios. You can also select one or more

system ETW Providers to create a unique provider configuration and resulting captured message set for a Live Trace Session; or you can select ETW Providers to add to the predefined configuration of a default **Trace Scenario** to enhance the scope of data capture. See Capturing Message Data for more information.

Moreover, the data selection aspect of the BSV model that extends to capturing data live enables you to choose specific data that you want your Live Trace Session to capture, for example, with the use of any of the following:

- **Fast Filter**
- **Keyword** bitmask filter
- **Session Filter**
- **EtherType** and **IP Protocol Number** filters in a remote **Trace Scenario**
- **Parsing Level** filter

You can also view data from a Live Trace Session in the exact same presentation formats that you can apply to data that you load into Message Analyzer through a Data Retrieval Session.

More Information
To learn more about applying filters to a Data Retrieval Session, see the following topic:
Selecting Data to Retrieve
To learn more about applying filters that are common to either session type, see the following topics:
Filtering Message Data
Setting the Session Parsing Level
To learn more about applying filters in a Live Trace Session, see the following topics:
PEF-NDIS Fast Filters
PEF-WFP Layer Set Filters
WebProxy Filters
Working with Session Filters in a Live Trace Session
Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog
Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog
System ETW Provider Event Keyword/Level Settings

Targeting Saved Data as an Input Source

2 minutes to read

From a Data Retrieval Session, you can locate saved data files and build a message collection that contains the data you want to target for input to Message Analyzer. The fastest way to access the configuration for a Data Retrieval Session is to click the **New Session** button on the Message Analyzer **Start** page to display the **New Session** dialog. To begin configuration for a Data Retrieval Session, click the **Files** button under **Add Data Source** in the **New Session** dialog. From the **Files** tab of the **New Session** dialog, you can target the files that contain the saved data you want to load into Message Analyzer. You can then customize your session configuration with the use of various features to optimize Message Analyzer performance in the data retrieval process.

More Information
To learn more about the input file types that Message Analyzer supports, see Locating Supported Input Data File Types.

Optimizing Data Retrieval

By selecting specific input files, specifying a **Session Filter**, and/or configuring a **Time Filter**, you can limit the data that is retrieved, so that smaller, more targeted message sets can be extracted from sources containing high volumes of data. This optimizes the data retrieval process because it results in the following:

- More manageable data sets that focus only on specific data, optionally in a configurable window of time.
- Better performance when loading message data.
- More effective data analysis.
- A focused set of messages that expedites the analysis process for others with whom you are sharing your results.

You can also optimize performance and target specific data to be retrieved by using the **Truncated Parsing** feature to detect input files that contain truncated messages and cause Message Analyzer to apply a pared-down parser set that reduces data loading time. Another feature that you can use to significantly improve performance is to set a **Parsing Level** to limit how far up the message stack Message Analyzer will parse.

Creating a Data Retrieval Session

For a general workflow that you can follow to create a Data Retrieval Session, see Configuring a Data Retrieval Session. This section also contains subtopics that describe the various features and functions that you can use to create and start a Data Retrieval Session.

Configuring a Data Retrieval Session

5 minutes to read

This section describes the typical tasks that you will perform when configuring a Data Retrieval Session, along with some background information on the features you will be using. Locating and selecting saved files or logs that contain the data you want to load into Message Analyzer is the only required task, while most others described in this section are optional, depending on what you wish to accomplish.

For example, if you are loading messages into Message Analyzer from a textual log file, you will need to make sure that you have a configuration file that enables Message Analyzer to parse such messages. You can do this by either selecting a built-in configuration file from the **Text Log Configuration** drop-down menu on the **Files** tab of the **New Session** dialog, or by custom-coding one, as described in Opening Text Log Files. On the other hand, specifying a **Session Filter** or a **Time Filter** is optional, although advised if you are working with large data files and you want to focus and limit data retrieval in a specific way.

In the discussions that follow, see the following figure for the location of referenced features.

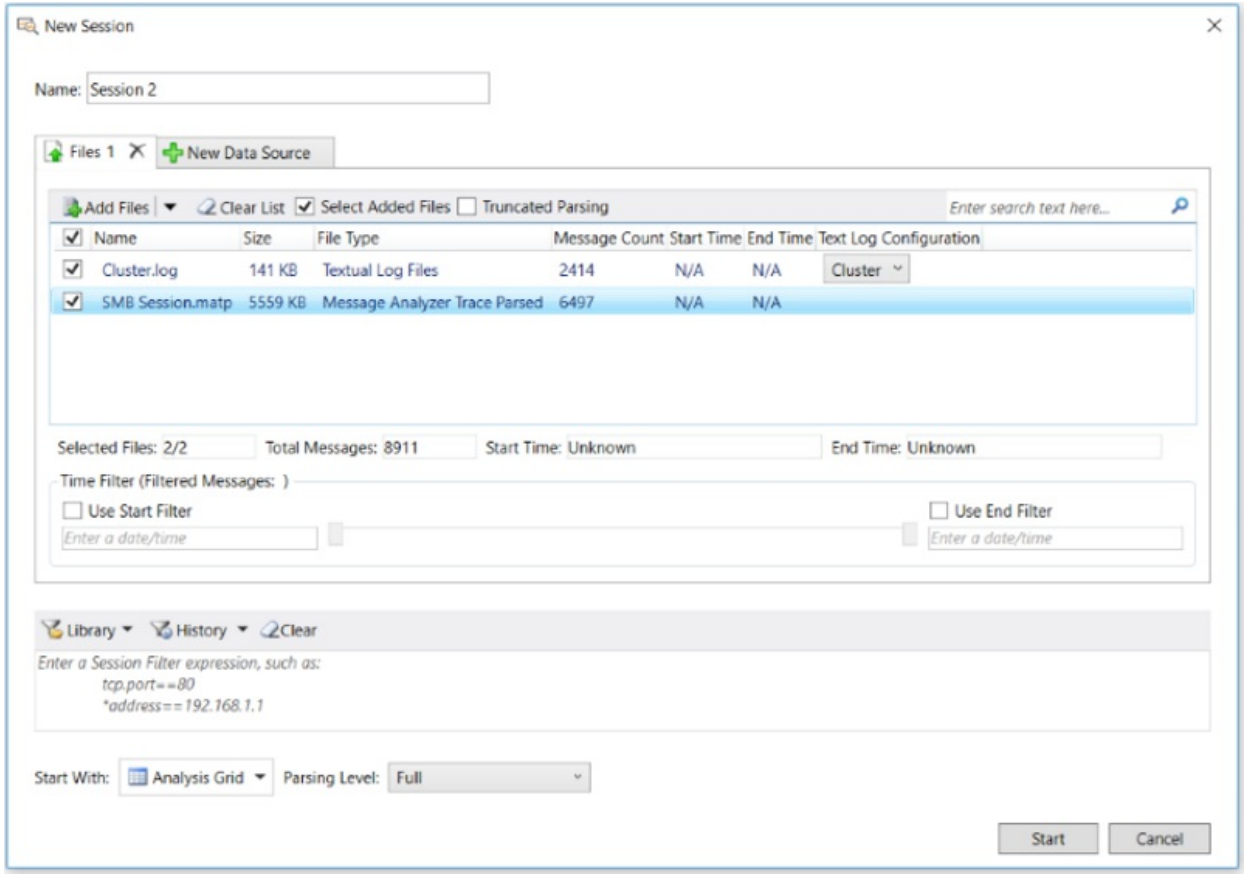

Figure 31: Data Retrieval Session configuration

Data Retrieval Session Workflow Overview

The following steps are an overview of the workflow that you can generally follow when configuring a Data Retrieval Session. Features for the following configuration tasks are accessible from the **Files** tab of the **New Session** dialog:

1. Verify that the input data files from which you will be retrieving data are file types that are supported by Message Analyzer, as described in Locating Supported Input Data File Types.

2. Target the message data to be retrieved from one or more data sources such as saved trace files or logs, as described in Performing Data Retrieval.

> **NOTE**
> You have the option to aggregate multiple data sources into a single session by making use of the **New Data Source** tab in a Data Retrieval Session. For additional details, see Configuring Session Scenarios with Selected Data Sources.

3. Select specific files that contain the data you want to work with, to create a subset of a larger targeted set of input files, as described in Performing Data Retrieval.
4. Optionally, if you have saved messages that are truncated, you can use the **Truncated Parsing** mode to handle trace files that contain such truncated messages, for example, a .cap file. This results in retrieving a smaller number of messages and improving performance, based on a pared-down message parser set, as described in Detecting and Supporting Message Truncation.
5. Specify a built-in or custom **Text Log Configuration** file that is required to parse a textual log file containing messages that you want to load and analyze with Message Analyzer, as described in Opening Text Log Files.
6. Optionally, if you have a very large set of input messages, you can configure and apply a **Time Filter** to create a precisely focused view of data in a specified window of time, as described in Applying an Input Time Filter to a Data Retrieval Session.

   > **NOTE**
   > Features for the configuration tasks that follow are accessible from outside the **Files** tab in the **New Session** dialog.

7. Optionally, configure and apply a **Session Filter** expression to the data being loaded to isolate specific data to be retrieved, as described in Applying a Session Filter to a Data Retrieval Session.
8. Optionally, choose a built-in **Parsing Level** scenario that limits the stack level to which Message Analyzer parses and provides filtering that creates a focused set of messages for analysis purposes. In addition, applying a **Parsing Level** can also dramatically improve performance, as described in Setting the Session Parsing Level.
9. Optionally, specify a data viewer in which to display the results of your Data Retrieval Session, other than the default viewer, as described in Selecting a Data Retrieval Session Viewer.
10. Optionally, specify a **Name** for your Data Retrieval Session and a **Description**, as described in Naming a Session.

Data Retrieval Session Filtering Overview

Two of the most important features that you can utilize to narrow the focus of data retrieval and significantly enhance Message Analyzer performance are the **Time Filter** and **Session Filter**. The following sections provide a brief overview of the advantages of using these filters in a Data Retrieval Session. Further details about these filters are described in the topics that are linked to in these sections.

Using a Time Filter

The Data Retrieval Session configuration options in the **New Session** dialog also include a **Time Filter** that enables you to select a time window in which to view data from files that are selected in the files list on the **Files** tab of the **New Session** dialog. This filter provides timeline slider controls with which you can set a time window. As you adjust these controls, Message Analyzer displays the time boundaries and the number of

messages contained in the window that you select, as described in Applying an Input Time Filter to a Data Retrieval Session.

This feature is useful when you have a large data set and you can estimate the time window in which a particular issue has occurred. By configuring a **Time Filter**, you can load, view, and analyze messages in a specified time window only, without incurring the additional overhead of loading the entire message set. If necessary, you can even edit the session and reconfigure the **Time Filter** so you can view messages in another time frame. However, **Time Filter** reconfiguration is available only in the **Full Edit** mode for which a button displays in the **Edit Session** dialog when you **Edit** an existing session, as described in Editing Existing Sessions.

> **NOTE**
> If you have an input file for which a Message Analyzer Data Retrieval Session does not display **Start Time** and **End Time** values, you can specify date-times in a format appropriate for the data file in the text boxes below the **Use Start Filter** and **Use End Filter** check boxes, as described in Applying an Input Time Filter to a Data Retrieval Session.

Using a Session Filter
If the input files for your Data Retrieval Session are large, you can limit the amount of data that you retrieve from such files and reduce consumption of system resources. You can do this by applying a **Session Filter** to the data to be loaded into Message Analyzer, to narrow the focus of retrieved data, improve performance, and streamline the analysis process, as described in Applying a Session Filter to a Data Retrieval Session.

Data Retrieval Session Configuration Features

The following subtopics describe the Data Retrieval Session configuration features that Message Analyzer provides and various operations that Message Analyzer supports:

Locating Supported Input Data File Types
Detecting and Supporting Message Truncation
Decrypting Input Data
Selecting Data to Retrieve
Selecting a Data Retrieval Session Viewer
Working With Special Input Requirements
Acquiring Data From Other Input Sources
Merging and Aggregating Message Data
Naming a Session

Retrieving the Data
When you are ready to load data into Message Analyzer, see Performing Data Retrieval to review various methods for retrieving saved data with Message Analyzer.

Locating Supported Input Data File Types

4 minutes to read

To locate saved data that you want to load into Message Analyzer from supported file types and to access session configuration features prior to loading the data, you should first create a Data Retrieval Session as described in Starting a Message Analyzer Session. To load data very quickly from supported file types into Message Analyzer without accessing the session configuration settings of the **New Session** dialog, you can use any of the methods described in Performing Data Retrieval. The exception to these "fast" methods is for text log files with a .log extension, where instead, Message Analyzer automatically opens the **New Session** dialog so that you can select a **Text Log Configuration** file that is required for parsing text logs.

> **NOTE**
>
> For Data Retrieval Sessions that contain a large message collection with many files targeted for loading data, you can search the files list to quickly locate any file by specifying file name characters in the search text box on the toolbar of the **Files** tab. For easy location capability, you should consider using relevant file naming conventions, as described in Naming Saved Files, when saving files from which you expect to reload data at some point.

Supported Input File Types

The data that you load into Message Analyzer must be derived from one or more of the supported file types, as described in the table that follows. Some of these file types are in Message Analyzer default native format while others are non-native. Some input files have certain requirements, as indicated in the table. Note that you can save Message Analyzer data in the .matp native file format only, even if you initially loaded data from a message file that has a non-native format.

Table 9. Message Analyzer Supported Input File Types

MESSAGE FILE NAME EXTENSION	DESCRIPTION	SUPPORT	USES/REQUIREMENTS
.aztable	Azure table storage	Non-native	Display data from Azure table storage (requires input connection information, as described in Retrieving Azure Storage Table Data).
.blg	Binary performance file (PerfMon)	Non-native	Standard input file support. A Message Analyzer **Profile** is available for this input file type.
.cap	Network Monitor capture file	Native	Open Network Monitor .cap files. Note that Message Analyzer can export trace data in the .cap file format. Several Message Analyzer **Profiles** are available for this input file type.

MESSAGE FILE NAME EXTENSION	DESCRIPTION	SUPPORT	USES/REQUIREMENTS
.csv	Comma Separated Value file	Non-native	Standard input file support.
.dmp	Crash Dump Files	Non-native	Standard input file support.
.etl	Windows event trace log file	Non-native	Open .etl files with embedded manifests containing event metadata per event. Otherwise, may require generating an ETL manifest; see Understanding Event Parsing with a Provider Manifest. Several Message Analyzer **Profiles** are available for this input file type.
.evtx	Event log file	Non-native	Standard input file support. A Message Analyzer **Profile** is available for this input file type.
.json	Javascript object notation log files	Non-native	Standard input file support.
.log	Textual log file	Non-native	Requires either selecting or creating an OPN Configuration file; see Opening Text Log Files. Note that the OPN configuration file selection requirement also applies to Azure storage logs. Message Analyzer **Profiles** are available for this input file type.
.matp	Compressed trace parsed file	Native	Open files saved in Message Analyzer native file format.
.matu	Compressed trace unparsed file	Native	Open Message Analyzer unparsed (.matu) files or PowerShell initiated traces that are saved in the same unparsed file format; see the Save-PefMessageCollection cmdlet documentation.
.oms	Operations Management Suite (OMS) logs	Non-native	Requires an Azure Subscription and credentials to log in to Azure, as described in Loading OMS Log Data.

MESSAGE FILE NAME EXTENSION	DESCRIPTION	SUPPORT	USES/REQUIREMENTS
.opn	Open Protocol Notation files	Native	Display the source code for any OPN message parser.
.pcap	Packet capture trace file	Non-native	Standard input file support. Message Analyzer **Profiles** are available for this input file type.
.pcapng	PCAP next generation trace file	Non-native	Standard input file support. Message Analyzer **Profiles** are available for this input file type.
.pmlcsv	Process monitor (procmon) trace file	Non-native	Requires changing a procmon file that was saved in .csv format to use the .pmlcsv file extension, to map properly to the correct Message Analyzer data loader.
.pmlxml	Process monitor (procmon) trace file	Non-native	Standard input file support.
.ps1	PowerShell script file	Non-native	Requires a script that runs a specified cmdlet, for example, to get event log data or target input files as a data source.
.saz	Session Archive Zip files	Non-native	Open Fiddler archives containing HTTP(s) data. A Message Analyzer **Profile** is available for this input file type.
.sqltable	SQL table files	Non-native	Standard input file support.
.trc	Network Associates Sniffer - DOS files	Non-native	Standard input file support.
.tsv	Tab Separated Value file	Non-native	Standard input file support.
.xml	XML file	Non-native	Standard input file support.

More Information

To learn more about how to use Message Analyzer **Profiles** to automatically display a predefined viewer and layout configuration for a targeted analysis context when loading data from applicable file types, see Working With Message Analyzer Profiles.

To learn more about different methods that you can use to quickly load data into Message Analyzer, see Performing Data Retrieval.

To learn more about saving your message data in the native .matp file format, see Saving Files in Native Format.

To learn more about how to create an OPN Configuration file for a text log, see the OPN Configuration File

for Textlog Adapter document download.

Detecting and Supporting Message Truncation

2 minutes to read

Truncation Support Message Analyzer provides support for input files in the .cap, .pcap, .pcapng, and .etl format that contain truncated messages. Message Analyzer can automatically detect whether messages in these input file types are truncated, at which time, it switches to a limited OPN parser set that contains parsers for the Ethernet, GRE, IPv4, IPv6, ESP, AH, IKE, AIPS, TCP, UDP, and HTTP protocols. Note that these particular parsers are specifically instrumented in Message Analyzer to support message truncation so that you can avoid unnecessary diagnosis messages in your loaded message results. When message truncation is detected, Message Analyzer also sets an **IsSessionTruncated** annotation that determines how individual messages will be parsed.

In addition, whenever Message Analyzer detects truncated messages, it automatically selects the **Truncated Parsing** check box on the toolbar of the **Files** tab in the **New Session** dialog, from where you locate and select files containing the data to load into Message Analyzer. Note that Message Analyzer will automatically select this check box whenever you have at least one input file that contains truncated messages. However, you have the option to manually unselect this check box. Moreover, if Message Analyzer does not detect message truncation in a particular input file and you know that it contains truncated messages, you can manually select the **Truncated Parsing** check box to force Message Analyzer to initiate the truncated parsing mode.

> **NOTE**
> When resaving your data from such a Data Retrieval Session, you can persist the truncation information when saving the data in the native Message Analyzer .matp file format only.

After you load files with truncated messages into the **Analysis Grid** viewer from a Data Retrieval Session, you can display the length to which messages in these files are truncated by adding a **TruncationLength** column to the **Analysis Grid** viewer from the **Field Chooser Tool Window** under the **Global Annotations** node. Thereafter, the **TruncationLength** column will be populated with truncation length values for the truncated messages.

Decrypting Input Data

2 minutes to read

Message Analyzer enables you to decrypt the messages in saved trace files from protocols that use the Transport Layer Security (TLS) and Secure Sockets Layer (SSL) encryption protocols, which includes HTTP and the Remote Desktop Protocol (RDP). However, to enable the Message Analyzer **Decryption** feature to decrypt messages in any particular saved trace, you will need to import an applicable server certificate into the Message Analyzer certificate store and specify a required password.

For further information about Message Analyzer decryption capabilities, which includes decrypting live traces, see Decrypting TLS and SSL Encrypted Data.

Selecting Data to Retrieve

2 minutes to read

As part of the BSV model, Message Analyzer provides a data selection feature that enables you to define the scope of the information that you load through a session. Although the concept of data selection applies equally to loading saved data in a Data Retrieval Session and capturing data in a Live Trace Session, this section focuses on selecting data in a Data Retrieval Session.

Using Data Selection

To use the data selection feature, you must configure your Data Retrieval Session to select specific data based on configurable criteria prior to starting the actual data retrieval process. During the retrieval process, the selection criteria that you configured is applied. This enables you to narrow the focus of the data retrieval process to only the message data that you want to work with. A very effective means of isolating different messages in a Data Retrieval Session, including those that have specific characteristics, consists of using a **Session Filter** to retrieve only the type of message data that you choose, while blocking all data that does not specifically meet your designated filtering criteria. Other aspects of data selection in a Data Retrieval Session consist of the following:

- Specifying the input file configuration with the use of check marks on the files in the files list of the **New Session** dialog, to select specific data from chosen files only.
- Adding a **Time Filter** that selects data according to the window of time you specify with the **Time Filter** controls in the **New Session** dialog.
- Choosing a **Parsing Level** that returns a set of messages that are constrained by the upper stack level to which Message Analyzer parses. This simultaneously creates a unique analysis perspective and focused message set, while improving performance by removing all messages above the specified **Parsing Level**.
- Enabling the **Truncated Parsing** mode, which enforces a limited parsing set to deal with data files that contain truncated messages, for example, a .cap file. This also improves performance and creates a unique analysis perspective by returning headers only for message types that are limited to a pared-down parser set, as described in Detecting and Supporting Message Truncation.

You might consider limiting data that you collect in a Data Retrieval Session by using a combination of a **Session Filter** and a **Time Filter**. For example, you could use a **Session Filter** to retrieve messages that transited a specified TCP port only and a **Time Filter** to limit the returned data to a particular window of time in which you suspect an issue is occurring with traffic on the specified port. The advantage of using a **Time Filter** with large data sets, such as you often have with log files, is that you can focus on specific subsets of data in order to reduce data loading time and to obtain better performance.

More Information

To learn more about how to select data from a Data Retrieval Session with a **Session Filter**, see Applying a Session Filter to a Data Retrieval Session.

To learn more about how to select data from a Data Retrieval Session with a **Time Filter**, see Applying an Input Time Filter to a Data Retrieval Session.

To learn more about how to choose the point in the stack up to which Message Analyzer will parse, see Specifying a Parsing Level.

See Also

Working with Session Filters in a Live Trace Session
Managing Session Filters

Applying a Session Filter to a Data Retrieval Session

2 minutes to read

After specifying one or more data sources in a Data Retrieval Session, you can apply a **Session Filter** expression to those data sources to filter the loaded message data to specific criteria. The same data selection capability is provided for Live Trace Sessions, where you can also apply a **Session Filter** expression to select specific data that is returned in your trace results. A **Session Filter** is shown in the topic Using a Session Filter.

To apply a **Session Filter**, you can either choose a built-in Filter Expression or you can create your own. When you create your own Filter Expression or modify a built-in filter, you should be careful to create a valid Filter Expression, or you may not return any data. For further information about creating valid Filter Expressions, see the topic Writing Filter Expressions to learn about the Filtering Language that is used by Message Analyzer.

To specify one of the built-in Filter expressions as a **Session Filter** in a Data Retrieval Session (or a Live Trace Session), you can select the filter you want to use from the centralized **Message Analyzer Filters** asset collection **Library** that is accessible on the toolbar above the **Session Filter** text box in the **New Session** dialog. To review further information on selecting a **Session Filter** in various categories that Message Analyzer provides, see the topic Selecting Built-In Session Filters. Note that the centralized Filter Expression **Library** is also accessible from the **Edit Session** dialog; and also from the Filtering toolbar just above the **Analysis Grid** viewer in any Analysis Session, where you can apply the same Filter Expressions to your trace results for data analysis purposes.

> **NOTE**
>
> If you have created and saved a Filter Expression that uses an **Alias** (uses a friendly name that replaces some cryptic field value) or a **Union** (correlates two data fields with similar values but different names in a single new field), that Filter Expression will appear in the centralized **Library** drop-down list wherever this **Library** is exposed in the Message Analyzer user interface (UI). You can use such a Filter Expression that contains an **Alias** or a **Union**, just as you would any other filter.

By using a **Session Filter** to select specific data from a Data Retrieval Session, you can work with specific data that you want to focus on, so you can get to the heart of the issue at hand without being overloaded with superfluous information.

> **NOTE**
>
> When you apply a **Session Filter** to a Data Retrieval Session, a funnel icon displays to the right of the corresponding top-level session node in the **Session Explorer Tool Window**, to indicate that the session has had a **Session Filter** applied to it. This icon provides a quick reminder of the analysis status of the session.

More Information

To learn more about **Session Filters**, see relevant information in the topic Working with Session Filters in a Live Trace Session.

To learn more about creating Filter Expressions and using the Message Analyzer Filtering Language, see Writing Filter Expressions.

To learn more about **Aliases**, see Using and Managing Message Analyzer Aliases.

To learn more about **Unions**, see Configuring and Managing Message Analyzer Unions.

Applying an Input Time Filter to a Data Retrieval Session

6 minutes to read

Prior to loading data into Message Analyzer from a specified message collection that you configure in a Data Retrieval Session, you can also configure a window of time in which to view data by using a **Time Filter**. This is particularly useful if you have one or more very large data files that contain trace or log data that was collected over a significant period of time and you want to narrow the scope of the data to be viewed. The UI configuration of the **Time Filter** feature is shown in the figure of the topic Configuring a Data Retrieval Session.

Most logs and file types are supported by the **Time Filter** feature; however, you may need to select the **Use Start Filter** and **Use End Filter** check boxes and specify start and end times in the corresponding text boxes for certain types of logs and traces in which the data format prevented Message Analyzer from retrieving this information. See Manually Specifying Time Formats for more information.

For example, the following table lists the different trace or log file types for which Message Analyzer can determine the start times and end times, along with those that it cannot.

Table 10. Estimated Start and End Times vs Trace File Types

TRACE TYPE	ESTIMATED START TIME	ESTIMATED END TIME	TOTAL MESSAGES/FILTERED MESSAGES
.aztable*	N/A	N/A	N/A
.blg	No	No	No
.cap**	**Yes**	**Yes**	**Yes**
.csv	No	No	No
.dmp	No	No	No
.etl	**Yes**	No	No
.evtx	No	No	No
.json	No	No	No
.log	**Yes**	**Yes**	**Yes**
.matp	**Yes**	**Yes**	**Yes**
.matu	**Yes**	**Yes**	**Yes**
.oms*	NA	NA	NA
.pcap	No	No	No

TRACE TYPE	ESTIMATED START TIME	ESTIMATED END TIME	TOTAL MESSAGES/FILTERED MESSAGES
.pcapng	No	No	No
.pmlcsv	No	No	No
.pmlxml	No	No	No
.ps1	No	No	No
.saz	No	No	**Yes**
.sqltable*	N/A	N/A	N/A
.trc	No	No	No
.tsv	No	No	No
.xml	No	No	No

Caution

*Message Analyzer supports these input file types with separate user interface (UI) features, as described in Retrieving Azure Storage Table Data, Loading SQL Data, and Loading OMS Log Data. However, the UI for these features does not provide start and end time information, as described in the previous table.

> **NOTE**
> **When loading data from a .cap file that was not previously saved by Message Analyzer, the start and end time values for the message data is unknown. In cases where Message Analyzer has *some* information, any displayed values for these files are only estimates. For text logs in .log file format, the start and end time values of messages will be known only after the parsing process, or if you specified the time stamp format in an OPN configuration file for the text log.

You should also be aware that estimated start and end times for some supported file types, as indicated in the **Start Time** and **End Time** text boxes in the **Time Filter** pane of your Data Retrieval Session configuration, may not match the **Start Time** and **End Time** text box values in the **Time Filter** panel on the Filtering toolbar in an Analysis Session. This is because the start and end times indicated in a Data Retrieval Session for certain file types is only an estimate, while the **Start Time** and **End Time** indicated in the **Time Filter** panel on the Filtering toolbar are known values that are determined in the parsing process. Moreover, this is the case because the start and end time estimations in the **Time Filter** pane are calculated *before* the data is retrieved, whereas the start and end time indications that appear on the Filtering toolbar are calculated *after* the data is retrieved.

Input File Details

After you add a supported input data file to your Data Retrieval Session and select it in the files list on the **Files** tab of the **New Session** dialog, Message Analyzer estimates the start and end times of the events in the trace or log files and displays these values, respectively, on each side of the adjustable time-window slider controls in the **Time Filter** pane of the Data Retrieval Session configuration. The start and end times are also displayed in the **Start Time** and **End Time** text boxes just above the **Time Filter** pane. In addition, an estimate of the total number of messages in the trace or log time frame is displayed in the **Total Messages** text box and the number of selected files in the files list is indicated in the **Selected Files** text box. Other file attribute information is also included, such as the file **Name**, **Size**, **File Type**, and **Message Count**; which all display in the columns just

below the toolbar on the **Files** tab of the **New Session** dialog.

> **NOTE**
> If you are loading multiple data sources in a Data Retrieval Session and any of those sources are among the ones listed in the previous table with a "No" in the estimated start time or estimated end time columns, the corresponding start or end times for the entire input file configuration will not display in the **Time Filter** section of the Data Retrieval Session configuration of the **New Session** dialog.

Configuring a Time Filter

To configure a **Time Filter**, you can adjust the time-window slider controls to zoom into a particular time slot in which you want to view data. Before you adjust the time slider controls, you should enable **Filtered Messages** tracking by selecting the **Use Start Filter** and **Use End Filter** check boxes. Thereafter, as you adjust the slider controls, a text box below the **Use Start Filter** check box displays the start time and a text box below the **Use End Filter** check box displays the end time in the **Time Filter** pane of the **New Session** dialog. The corresponding new start and end time values define the selected time window in which you are choosing to view data, while the parenthetical **Filtered Messages** label displays the estimated number of messages that are contained in the time slot that you specify. However, the overall start and end times of the trace or log persist in the **Start Time** and **End Time** text boxes above the **Time Filter** pane.

> **NOTE**
> Selecting data with a **Time Filter** produces a subset of the total messages contained in the selected input file/s. Once you apply a **Time Filter** that reduces message count in your Data Retrieval Session results, you will need to remove the **Time Filter** and rerun the Data Retrieval Session if you want to display the original message set in its entirety.
>
> To do this or to specify a different window of time, open the **Edit Session** dialog either by clicking the **Edit Session** icon on the global Message Analyzer toolbar or by selecting the **Edit Session** item from the Message Analyzer **Session** menu. After you open the **Edit Session** dialog, you will need to click the **Full Edit** button to enable the **Time Filter** controls. From the **Edit Session** dialog, you can then remove the previous **Time Filter** configuration altogether by moving the time slider controls all the way to the left and right side default positions, or you can reconfigure the original Data Retrieval Session with a new **Time Filter**. To apply the changes that you make to the **Time Filter** configuration, click **Apply** in the **Edit Session** dialog.

Manually Specifying Time Formats

To enable manual entry of time window boundaries in a specific format, select the **Use Start Filter** and **Use End Filter** check boxes. These selections make the corresponding time boundary text boxes writeable, so that you can specify time values in a format that is suitable for the displayed data. Thereafter, as you adjust the time slider controls, the changing time window boundary values match the time stamp format of the displayed data. This feature accommodates for message data that may have time stamps in a format that Message Analyzer cannot adequately determine, as indicated in the previous table.

The **Time Filter** configuration that you specify is then applied by Message Analyzer when you start the Data Retrieval Session with a click the **Start** button of the **New Session** dialog.

More Information
To learn more about reconfiguring a Data Retrieval Session and re-running it, see Editing Existing Sessions.
To learn more about creating an OPN configuration file for text logs with a proprietary format, the OPN Configuration File for Text Log Adapter document is available as a TechNet download.

Specifying a Parsing Level

2 minutes to read

If you are loading large traces into Message Analyzer, you might consider setting a built-in **Parsing Level** to improve performance. Because **Parsing Level** scenarios typically limit how far up the message stack Message Analyzer will parse, you can reduce the overall number of messages processed and dramatically improve performance. Also, by choosing a **Parsing Level**, you can create a more targeted message set to improve your analytical focus. This is because most **Parsing Level** scenarios contain predefined filters that are beneficial to message analysis at the top stack level on which any particular **Parsing Level** scenario creates focus.

More Information

To learn more about **Parsing Levels** and the built-in scenarios that you can employ, see Setting the Session Parsing Level.

Selecting a Data Retrieval Session Viewer

3 minutes to read

After you configure a Data Retrieval Session with the files that contain the data you are targeting to load into Message Analyzer, you can choose a viewer in which to display the session results data. You can choose any data viewer that exists in the **Start With** drop-down in the **New Session** dialog for a Data Retrieval Session, including data viewer assets that you import through the Message Analyzer sharing infrastructure. For a list of the data viewers that Message Analyzer provides by default, see the topic Selecting a Session Data Viewer.

Changing Data Viewers

You can start loading data into Message Analyzer through a Data Retrieval Session any time after you create a selected input file configuration in the files list on the **Files** tab of the **New Session** dialog. Unless you specify otherwise, Message Analyzer will use the default data viewer to display your data, which is typically the **Analysis Grid** viewer. However, you can change the default data viewer setting that will be used by Message Analyzer for all sessions by specifying a viewer in the **Default Viewer** drop-down list on the **Default Profile** pane of the global **Options** dialog that is accessible from the Message Analyzer **Tools** menu.

Note that you have the option to override the default data viewer prior to starting a session, by explicitly specifying a different one that might have certain properties that are more useful for the data you are loading and/or the type of analysis you are planning to do. You can do this by clicking the **Start With** drop-down list in the **New Session** dialog and then selecting a data viewer of choice.

Specifying a Default Data Viewer
The default data viewers that are available from the **Start With** drop-down list during Data Retrieval Session configuration, consist of the following:

- **Analysis Grid**
- **Grouping**
- **Pattern Match**
- **Gantt**
- **Chart** — displays a default Bar element graph with message volumes per module only when chosen from the **Start With** drop-down.

Note that the data viewers in the following list are Preview features. To make these viewers available for selection during session configuration, you must select them on the **Features** tab of the **Options** dialog, which is accessible from the Message Analyzer **Tools** menu, and then you will need to restart Message Analyzer.

- **Interaction**
- **Message Summary Tiles**
- **Message Summary Lists**
- **Perfmon**

Selecting Data Viewers After Data Retrieval
After you retrieve your data and it displays in the viewer that you initially selected, you can present your data in any of the other available viewers by selecting them from a context menu that is accessible by right-clicking any session node in the **Session Explorer Tool Window**. You can also access these same viewers by clicking the

New Viewer drop-down list on the global Message Analyzer toolbar. By selecting different data viewer configurations, you can achieve unique data analysis perspectives to assist in problem solving.

Note that for many of the data viewers that are available from the **New Viewer** drop-down list, a submenu displays to enable you to select a **Layout** for a chosen viewer. For example, this could be a **Layout** for the **Analysis Grid**, **Chart**, or **Grouping** viewer that each expose different message fields to create a focused analysis environment. Some examples of analysis environments that view **Layouts** can create include the following:

- SMB file sharing and performance
- Event log and Cluster log analysis
- HTTP Operations response time analysis
- Traffic volume distributions per module or conversation
- TCP diagnosis
- IP conversation bandwidth consumption

More Information
To learn more about the data viewers and **Layouts** that are available in Message Analyzer, along with the **Tool Windows** with which they interact, see the Data Viewers and Tool Windows sections of this Operating Guide.
To learn more about the integrated and interactive analysis environments that Message Analyzer provides with preset configurations of data viewers, **Layouts**, and **Tool Windows**, see Working With Message Analyzer Profiles.

See Also

Session Data Viewer Options

Working With Special Input Requirements

2 minutes to read

This section describes how you can handle input data that has a unique format, in addition to cases where the data you are loading was captured with a provider for which your system has no provider manifest.

Handling Unique Log File Formats

A common log file type that Message Analyzer supports uses the .log extension, which can have a proprietary format. Prior to loading data into a Message Analyzer Data Retrieval Session from such a text-based log file, you will typically need to select a predefined parser from the **Text Log Configuration** drop-down list above the files list in your Data Retrieval Session. If you do not select one of these predefined OPN configuration files, or if an appropriate configuration file is unavailable, it is likely that Message Analyzer will not recognize the format of your log and will therefore be unable to parse it fully. If this occurs, Message Analyzer will display all of the message fields from your log in the **Summary** column of the **Analysis Grid** viewer with each field delimited by a semicolon.

If you want Message Analyzer to successfully parse and display all the message fields of your log in separate **Analysis Grid** columns, you will need to create a custom OPN configuration file to parse the log. Moreover, if Message Analyzer does not have an OPN description for any particular message that is contained in a particular log file from which you are trying to load data, you will be unable to display any data for *that message*. This is because, in absence of an OPN description, the protocol object model (POM) (see PEF Architecture Tutorial) will not contain a *compiled* OPN description that the PEF Runtime can use to parse the messages. When you create an OPN configuration file, you will define the log messages with OPN declarations and Regex notation, as described in the OPN Configuration File for Text Log Adapter v2 document. When complete, you must drop the new configuration file into the following directory location:

```
%LocalAppData%\Microsoft\MessageAnalyzer\OpnAndConfiguration\TextLogConfiguration\DevicesAndLogs\
```

Thereafter, whenever you add a log file of this particular type to the files list in a Data Retrieval Session, an OPN description will be automatically generated from the content of the existing OPN configuration file after you select the file from the **Text Log Configuration** drop-down list in session configuration and you **Start** the session.

Missing Provider Manifests

Message Analyzer also accommodates for loading messages that were captured with a provider for which your system has no provider manifest, if the proper manifest is included with the loaded trace file. In addition, if you have a destination computer with a manifest that is out-of-date with respect to the source computer manifest that was included in an .etl file, a workaround that you can employ to create the manifest that you require is described in Generating a Provider Manifest. Otherwise, Message Analyzer will be unable to parse messages from such message providers.

> **NOTE**
>
> Message Analyzer saves the OPN parser configuration with all files that you save in .matp format, to make the data portable from one Message Analyzer installation to another should any differences exist in the parser packages.

More Information

To learn more about working with text-based log files, see Opening Text Log Files.

To learn more about creating an OPN configuration file for text-based logs, download the OPN Configuration File for Text Log Adapter v2 document.

Opening Text Log Files

4 minutes to read

Message Analyzer enables you to parse data from text-based log files, for example, files with a *.log extension. You can open these types of log files from the **New Session** dialog in a Data Retrieval Session, or by using **Windows Explorer**, the **Open** feature, or the drag-and-drop method.

> **NOTE**
>
> The drag-and-drop method is unavailable when running Message Analyzer in the Administrator mode due to varying security contexts.

Using a Configuration File to Parse Text Logs

To parse text-based log files, Message Analyzer requires a configuration file. By default, Message Analyzer provides several predefined configuration files that you can select from a drop-down list in the **Text Log Configuration** column just below the toolbar on the **Files** tab of the **New Session** dialog during Data Retrieval Session configuration. You can view the **Text Log Configuration** drop-down list in the figure of the topic Parsing Input Text Log Files.

If there is no predefined configuration file that meets your requirements, you are advised to create an OPN configuration file so that Message Analyzer can fully parse your text log. Thereafter, you will need to either set the new configuration file as the default for all your log files, by selecting it from the drop-down list in the **Text Log Files** section on the **General** tab of the global **Options** dialog — which is accessible from the Message Analyzer **Tools** menu — or you must select a unique configuration file in the **Text Log Configuration** drop-down for each specific log file that contains a varying message format. For more information about configuration tasks that are required for opening and parsing text logs, see the Addendum 1: Configuration Requirements for Parsing Custom Text Logs topic.

After you create a configuration file for your text log and place it in the appropriate directory location, as indicated in the Addendum 1: Configuration Requirements for Parsing Custom Text Logs topic, you can then target the text log that contains the messages to load into Message Analyzer through the **Add Files** feature, which is located in the toolbar on the **Files** tab of the **New Session** dialog. After you select a log file with the **Add Files** feature and it displays in the files list, the **Text Log Configuration** drop-down menu is populated with the names of all the predefined configuration files that ship with Message Analyzer, in addition to any that you have created. At this point, you can simply select the appropriate configuration file and start the parsing process by clicking the **Start** button in the **New Session** dialog, at which time messages begin loading into the specified data viewer, for example the **Analysis Grid**.

Caution

One of the selections available in the drop-down list under the **Text Log Configuration** column for each text log is the **[None]** option. If you select this option and start to load data, Message Analyzer does not fully parse the messages in the log, but instead simply returns the lines of message data in the file. In this case, it is likely you will see data only in the **MessageNumber**, **Module**, and **Summary** columns of the **Analysis Grid** viewer, where the log data displayed in the **Summary** column contains message fields that are delimited by semicolons. Also, if you select an invalid configuration file for your log, Message Analyzer will display a diagnostic **Parsing Error** icon in the **DiagnosisTypes** column of the **Analysis Grid** viewer for each log message.

TIP

If you do not create a configuration file for your text log and you simply select the **[None]** option in the **Text Log Configuration** column drop-down list, you can still use Message Analyzer filtering features to locate specific data from your log file after the data displays in the **Analysis Grid** viewer. For example, you might specify a view **Filter** such as `contains “error”` or `*Summary contains “error”` to filter for lines of data that contain “error” characters, or you could also specify **Color Rule** filters or a **Pattern Expression** to highlight or extract other data of interest.

Using Alternate Input Methods for Text Log Files

You have the option to load data from a .log file by using **Windows Explorer**, the **Open** feature, or the drag-and-drop method. Providing that you have not set a default configuration file, Message Analyzer automatically opens to the **New Session** dialog to enable you to select a configuration file from the **Text Log Configuration** drop-down menu on the **Files** tab. However, if you have already set a default configuration file that your text logs will use (from the **General** tab of the global **Options** dialog that is accessible from the global Message Analyzer **Tools** menu), as described in the Addendum 1: Configuration Requirements for Parsing Custom Text Logs, then Message Analyzer automatically parses and displays the data in the default data viewer without displaying any Data Retrieval Session configuration options.

To open multiple log files simultaneously from outside a Data Retrieval Session, you can double-click supported files in **Windows Explorer**, use the drag-and-drop method, or use the **Open** feature. Note that if you use **Windows Explorer**, you might need to select the right-click **Open With** context menu command to associate Message Analyzer as the application that opens your log files.

NOTE

You can also use **Windows Explorer**, the **Open** feature, or the drag-and-drop method as alternate methods of loading data from other supported file types, as described in Locating Supported Input Data File Types.

Acquiring Data From Other Input Sources

3 minutes to read

Beyond live captures and saved files, Message Analyzer enables you to utilize additional sources to acquire input data for enhanced analysis capabilities. As provided in the **New Session** dialog, you can create the configuration for a Live Trace Session by clicking the **Live Trace** button under **Add Data Source** to capture data live from the network. You can also create the configuration for a Data Retrieval Session by clicking the **Files** button under **Add Data Source** in the **New Session** dialog, which enables you to target save files (traces and logs) as input to Message Analyzer. However, there are several other input sources under **Add Data Source** from which you can load data into Message Analyzer through a Data Retrieval Session, as follows:

> **NOTE**
>
> WPP-generated events are included in the list below because they are an additional input data source for Message Analyzer. However, the **Add Data Source** feature in the **New Session** dialog does not provide access to the configuration required to set up for processing WPP events. To learn more about how to capture or retrieve WPP events, see Loading WPP-Generated Events.

- **Azure data** — provides access to data from an Azure table. You can create the input configuration by clicking the **Azure Table** button under **Add Data Source** in the **New Session** dialog.

You an also acquire input data from Azure storage binary large object (BLOB) logs with the use of the **File Selector**, which is accessible from the global Message Analyzer **File** menu, by selecting **Open** and then clicking the **From Other File Sources** command.

- **Event logs** — a preview feature that provides access to data from **Microsoft Event Viewer** logs, such as **Applications and Services**, **Windows**, and others. You can create the input configuration by clicking the **Event Logs** button under **Add Data Source** in the **New Session** dialog.
- **PowerShell query** — a preview feature that provides access to data that is output by a PowerShell query, which you create with one or more PowerShell cmdlets. You can create the input configuration by clicking the **PowerShell** button under **Add Data Source** in the **New Session** dialog.
- **SQL query** — a preview feature that provides access to data from a SQL database table. You can create the input configuration by clicking the **Sql** button under **Add Data Source** in the **New Session** dialog.
- **WPP-Generated Events** — Message Analyzer can process Windows software trace preprocessor (WPP)-generated events. Because WPP events make use of the ETW framework, Message Analyzer can capture them live or load them from a saved event trace log (ETL) file.
- **Operations Management Suite (OMS) logs** — Message Analyzer can load OMS log data, which enables you to leverage Message Analyzer data viewers and analysis capabilities when working with this type of data. To facilitate the process, Message Analyzer provides a search interface to OMS Log Analytics that you can access through the **Oms** data source feature of the **New Session** dialog during Data Retrieval Session configuration. Note that OMS is a preview feature.

> **IMPORTANT**
>
> To use a preview feature, ensure that it is selected in the **Preview Features** list on the **Features** tab of the **Options** dialog, which you can access from the global Message Analyzer **Tools** menu. If you enable a previously disabled feature, you will need to restart Message Analyzer in order to use the feature.

The above input data sources are described in the following topics of this section:

Handling Azure Data Loading System Event Log Data Deriving Input Data with PowerShell Scripts Loading SQL Data Loading WPP-Generated Events Loading OMS Log Data

Combining Input Data Sources

You can use any of the previously specified input data sources alone, or you can combine two or more input sources so that you can correlate and analyze related data from different sources. For example, you might combine a saved trace file with a particular Event Log to correlate data from captured messages versus event log entries, for enhanced analysis. Given that some of the input data sources are preview features, you might also come up with some inventive ways to combine data sources in other ways. For example, you might combine a PowerShell query or a SQL database table with another input source. In the case of a PowerShell query, you could acquire process IDs on a specified remote computer and correlate that data with a network trace saved in an ETL file. For this particular scenario, Message Analyzer has some assets that are specifically designed to process this type of data, for example, the **Grouping** viewer which can correlate process names and IDs.

See Also

Configuring Session Scenarios with Selected Data Sources Grouping Viewer

Handling Azure Data

2 minutes to read

To enhance your troubleshooting experience in the Cloud environment, Message Analyzer provides a simple way for you to browse for, select, and view data from log files that are stored in Azure binary large object (BLOB) containers. To access Azure logs in these containers, Message Analyzer provides a special dialog called the **File Selector**, which enables you to specify Azure connection information, navigate through folders and files in the BLOB containers, and then select the Azure log files you want to open. Message Analyzer also enables you to load input data from Azure tables by creating an **Azure** input configuration from the **New Session** dialog that specifies Azure connection information. The primary difference between Azure table data and the data obtained from an Azure log that is stored in a BLOB container is the format. Azure tables store messages in each row, whereas each column contains a corresponding message property that Message Analyzer parses as a field. Azure log data is typically stored in .log file format, which requires an OPN configuration file for Message Analyzer to parse it, as described in Using the AzureStorageLog Parser.

> **NOTE**
>
> The **File Selector** was created to access stored files that Windows **File Explorer** does not support.

Limiting the Azure Data Retrieved

Just like any other input source that can provide data to Message Analyzer, you can configure and apply a **Time Filter** to limit the amount of Azure data you are loading, for better performance and to create focus on specific data that you want to view. Also, in the case of Azure log files, you can include a **Session Filter** from the **Azure Storage** category in the Message Analyzer Filter Expression **Library** that is accessible in the **New Session** dialog, to further narrow the scope of the data you want to retrieve. Because you can select one or more log files from Azure Storage BLOB containers as input to Message Analyzer, these filtering techniques are essential to facilitate cross-log analysis when you are merging and aggregating similar data from multiple high-volume logs. However, note that you might need to download the **Azure Storage Filters** asset collection with the use of the **Asset Manager** dialog, as accessible from the Message Analyzer **Tools** menu, to update your centralized Filter Expression **Library** to contain appropriate Azure filters in the **Azure Storage** category.

Viewing Azure Data

After Message Analyzer loads the data from your Azure log/s or table, you can view it as fields of data in the **Analysis Grid** viewer columns. You can also employ the **Field Chooser Tool Window** to add new data columns to the **Analysis Grid** viewer based on field names that are specific to the Azure storage logs or table, for enhanced analysis.

Specifying Azure Connection Information

In order to access Azure data from the previously mentioned sources, you will need to provide specific connection information. In the case of Azure tables, you will also need to enable the **Azure Table Import** preview feature on the **Features** tab of the global **Options** dialog and restart Message Analyzer.

Loading Azure Data

The following subtopics of this section describe how to get started with loading Azure data into Message Analyzer.

Retrieving Azure Storage Blob Data
Retrieving Azure Storage Table Data

Go To Procedure

For an example of how to retrieve Azure data, see the procedural topic Retrieve Data from Log Files in Azure Storage BLOB Containers or from an Azure Table.

Retrieving Azure Storage Blob Data

5 minutes to read

This topic shows you how to access and load data into Message Analyzer version 1.3 and later from log files that exist in Azure Storage binary large object (BLOB) containers. For Message Analyzer to successfully load and display this log data, you will need to provide some input connection information in order to access the BLOB repositories where the log data is saved. To access these repositories, you will use the **File Selector** dialog to select the files that contain the data you want to work with. After you specify such files and you exit the dialog, you must select a particular configuration file in your Data Retrieval Session for parsing the log data. If the **AzureStorageLog** configuration file does not exist in the **Text Log Configuration** drop-down list in your Data Retrieval Session configuration, you will need to download it through the Message Analyzer sharing infrastructure, as described in Using the AzureStorageLog Parser, and then restart Message Analyzer.

Go to Procedure

To proceed directly to a detailed procedure that you can follow to retrieve log data from Azure Storage BLOB containers, see To access, load, and view log data from Azure storage BLOB containers. Otherwise, you can review the Workflow Overview.

Accessing Log Data in Azure Storage BLOB Containers

To connect with the Azure Storage BLOB repositories, you will need an **Account name** and **Account key**, which you should be able to obtain from your Azure portal. To input this information, you will need to open the **File Selector** dialog, which is accessible by clicking **Open** from the Message Analyzer **File** menu and then selecting **From Other File Sources**. Thereafter, the **File Selector** dialog enables you to enter the connection information in the **Add Azure Storage Connection** dialog that displays when you click the **Add Azure Connection** button. After you specify the input connection data, you can navigate to the Azure Storage BLOB containers and select one or more log files from which to load data into Message Analyzer.

> **NOTE**
>
> It is necessary to specify the input connection information only once for access to Azure logs, as long as you are still running the same Message Analyzer session where you initially specified such connection information. Otherwise, for security purposes, Azure account access is not persisted across Message Analyzer sessions.

The user interface for the **File Selector** dialog is shown in the figure that follows.

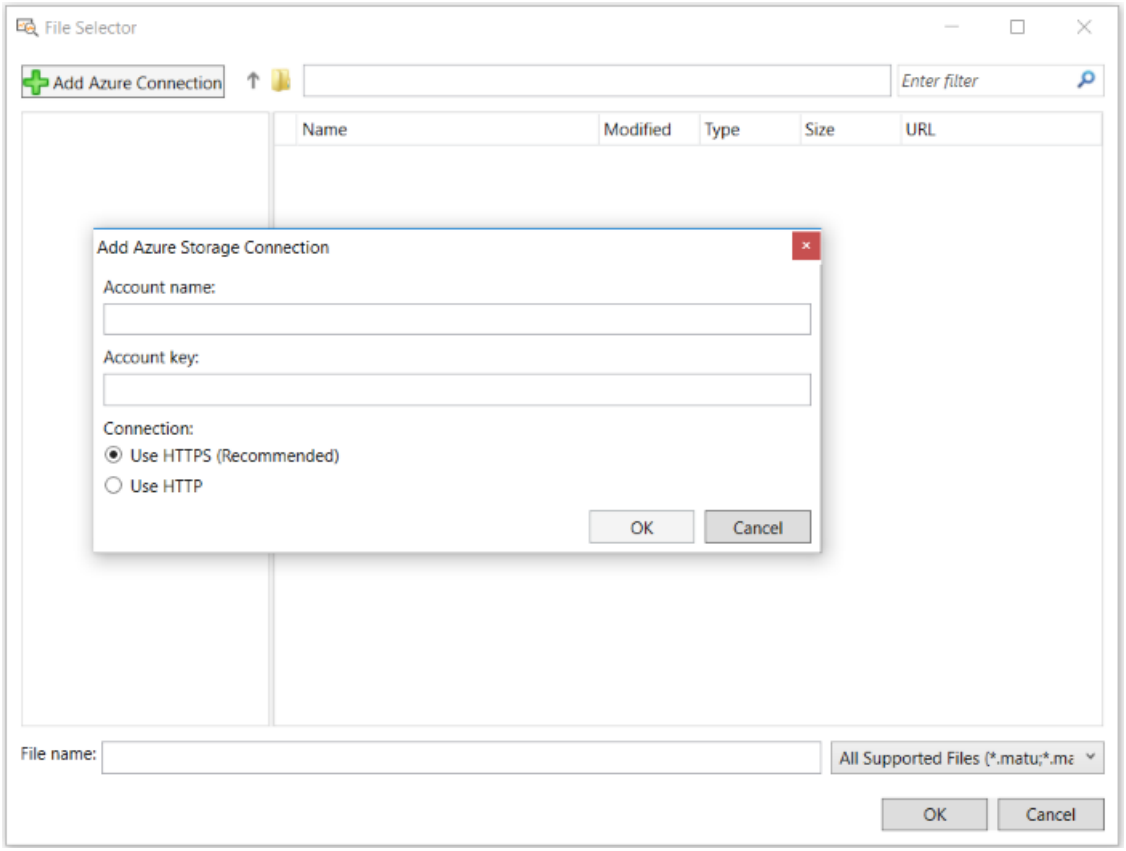

Figure 32: Azure File Selector dialog

Using the AzureStorageLog Parser

To load log file data into Message Analyzer from Azure Storage BLOB containers, a **Text Log Configuration** file is required for parsing the log data. After you target specific log files for input data and you click **OK** to exit the **File Selector** dialog, Message Analyzer displays the **New Session** dialog for a Data Retrieval Session, where the logs you specified are listed on the **Files** tab. In the session configuration, you will need to select the **AzureStorageLog** configuration file in the **Text Log Configuration** drop-down list to parse your Azure logs. If this file does not exist in the list, you will need to download it through the Message Analyzer sharing infrastructure.

The Azure parsers are available in the **Azure Storage Parsers Version x.x** asset collection that is accessible by clicking the Message Analyzer **Tools** menu and then selecting the **Asset Manager** command to display the **Asset Manager** dialog. To acquire the parsers in this collection, click the download icon (with the down-arrow indicator) to the right of the **Azure Storage Parsers** package on the **Downloads** tab of the **Asset Manager** dialog and then select the auto-sync option that appears in the **Item Download Options** dialog. When you click **OK** to exit this dialog, the Azure parser collection is downloaded to your Message Analyzer installation. Thereafter, you will receive automatic collection updates as they become available, as described in Downloading Assets and Auto-Syncing Updates.

You can now select the **AzureStorageLog** configuration file in the **Text Log Configuration** drop-down list during Data Retrieval Session configuration.

Analyzing Azure Data

After Message Analyzer loads your Azure data, you can analyze and correlate the fields from the parsed input log file data by viewing it in **Analysis Grid** viewer columns. You can also make use of the **Field Chooser Tool Window** to display additional Azure data fields of interest, by adding them as columns to the **Analysis Grid** viewer for further analysis. The additional fields appear under the **AzureStorageLog** node when it is expanded in the **Field Chooser** window.

Workflow Overview

The following procedure outlines the general steps you can follow when configuring a Data Retrieval Session to load log data from Azure Storage BLOB containers into Message Analyzer:

1. Ensure that you have the **AzureStorageLog** parser before you get started. If you need this parser, follow the general instructions in Using the AzureStorageLog Parser to obtain it.
2. Open the **File Selector** dialog from the Message Analyzer **File** menu by highlighting **Open** and then selecting the **From Other File Sources** command.
3. In the **Add Azure Storage Connection** dialog, specify **Account name** and **Account key** information, as described in Accessing Log Data in Azure Storage BLOB Containers.
4. In the **Add Azure Storage Connection** dialog, ensure that the **Use HTTPS (Recommended)** option is selected as the **Connection** protocol.
5. Expand the nodes in the **File Selection** dialog until you expose the **Blobs** container and subfolders, so you can navigate to the log data you want to retrieve.
6. Select one or more .log files and then click **OK** to exit the **File Selector** dialog, at which time the **New Session** dialog displays to enable additional configuration of your Data Retrieval Session.
7. Select the **AzureStorageLog** configuration file in the **Text Log Configuration** drop-down list for each Azure log file in the files list of the **New Session** dialog.
8. Optionally, configure a **Time Filter** in the **New Session** dialog to narrow the scope of data retrieval, as described in Applying an Input Time Filter to a Data Retrieval Session.
9. Optionally, if you have the **Azure Storage Filter** asset package downloaded, select a predefined Azure filter from the **Azure Storage** category in the Filter Expression **Library** to create a resulting data set that focuses on specific information, as described in Applying a Session Filter to a Data Retrieval Session.
10. Ensure that the **Analysis Grid** viewer is selected in the **Start With** drop-down list and click **Start** to exit the **New Session** dialog and begin data retrieval.
11. Observe that the log data displays in the **Analysis Grid** viewer.
12. Optionally, use the **Field Chooser** window to enhance the view layout for the log data by adding one or more **Analysis Grid** viewer columns based on other fields of the parsed Azure log data.

    Optionally, if you have the **Azure Storage View Layouts** asset package downloaded, change the **Analysis Grid** viewer column layout by selecting the **Storage Log** item in the **Layout** drop-down list. This layout is specifically designed to include key data fields for analysis of Azure storage logs.

More Information

To learn more about downloading and auto-syncing Message Analyzer assets, see Managing Asset Collection Downloads and Updates.

To learn more about the **Field Chooser**, see the Field Chooser Tool Window topic.

See Also

Retrieving Azure Storage Table Data

Retrieving Azure Storage Table Data

3 minutes to read

Message Analyzer versions 1.2 and later provide a data input preview feature known as the **Azure Table Import**, which enables you to retrieve data stored in an Azure Storage table. To enable this feature, you will need to select it on the **Features** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. After you click **OK** to exit the **Options** dialog, you must restart Message Analyzer so that the **Azure Table** button appears in the **New Session** dialog, from where you can select it.

Go To Procedure
To proceed directly to a detailed procedure that you can follow to retrieve Azure Storage table data, see To access, load, and view data stored in an Azure table. Otherwise, you can review the Workflow Overview.

Specifying Connection Information

After you start a **New Session** from the Message Analyzer **File** menu and you click the **Azure Table** button, an input configuration appears on the **Azure Table** tab from where you can provide connection information. To enable Message Analyzer to read Azure Storage table data, you will need to obtain and enter the following required information:

- **Account Name** — an Azure Storage account name. You can obtain this information from your Azure portal.
- **Account Key** — an Azure Storage account access key. You can obtain this information from your Azure portal.
- **Table Name** — the name of the table that contains the data you want to access.

The figure that follows shows the **Azure Table** interface from where you provide the specified input information.

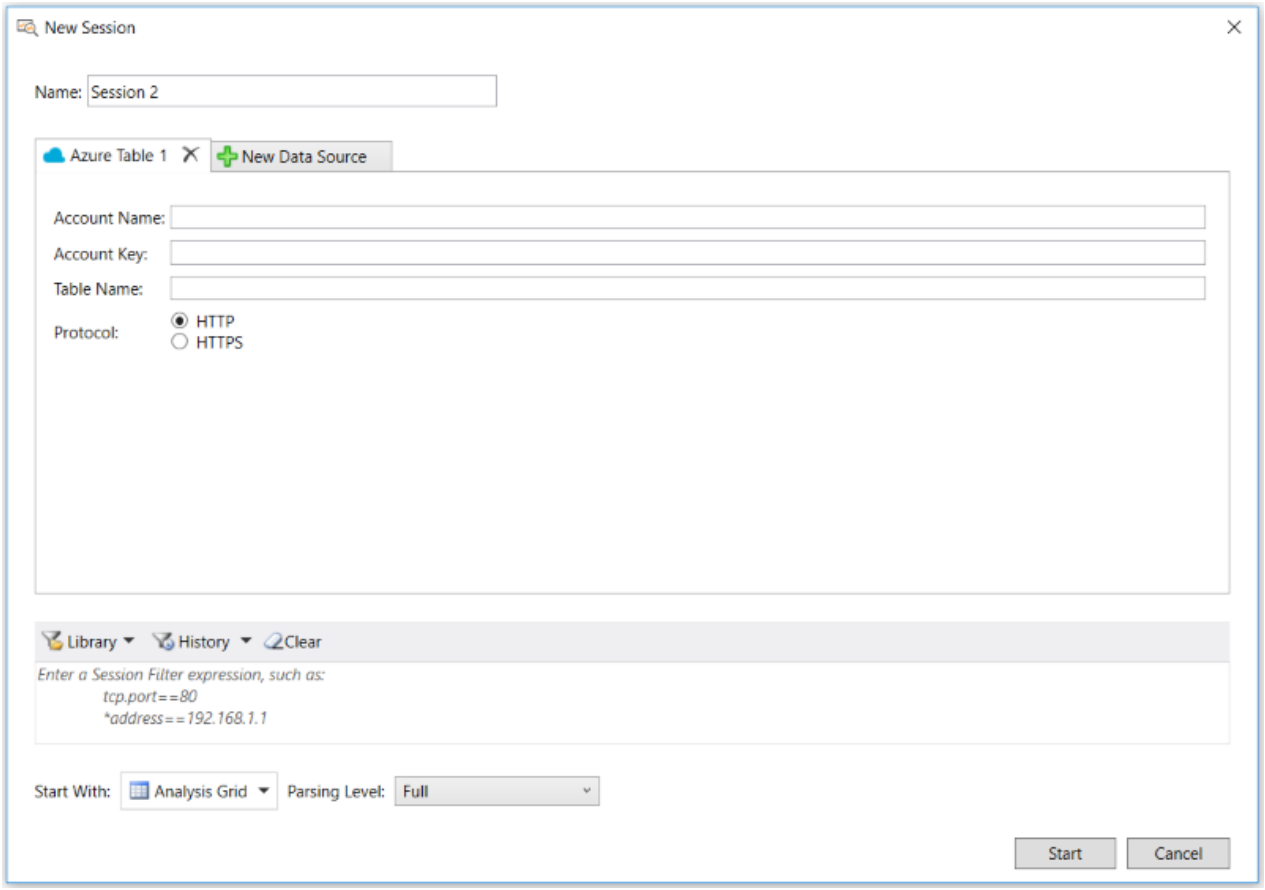

Figure 33: Azure Table input configuration

Workflow Overview

The following procedure outlines the general steps you can follow when configuring a Data Retrieval Session to load data from an Azure Storage table into Message Analyzer:

1. On the Message Analyzer **File** menu, point to **New Session** and then select **Azure Table** in the submenu to open a Data Retrieval Session that enables you to target Azure Storage table data as input to Message Analyzer.
2. Enter an **Account Name**, **Account Key**, and **Table Name** on the **Azure Table** tab of the **New Session** dialog.
3. Select either **HTTP** or **HTTPS** as the connection **Protocol**.
4. Ensure that the **Analysis Grid** viewer is selected in the **Start With** drop-down list.
5. Start retrieving data by clicking the **Start** button in the **New Session** dialog.
6. Observe that the Azure Storage table data displays in the **Analysis Grid** viewer, with the Azure table properties displaying as fields in the **Summary** column.
7. Optionally, expose other Azure table fields in the **Analysis Grid** viewer for enhanced analysis, by using the **Field Chooser Tool Window**.

NOTE

By default, Message Analyzer displays Azure Storage table information as rows of message data with all the parsed properties displaying in the **Summary** column of the **Analysis Grid** viewer. Because Message Analyzer parses each property in an Azure Storage table and represents it as a field after loading the data, you can utilize the **Field Chooser** window to configure separate columns in the **Analysis Grid** viewer to contain the field data for any of the Azure table properties that are normally condensed into the **Summary** column. This can make it easier to review and work with the values of individual fields.

You can locate the Azure Storage table properties as fields in the **Field Chooser** window under the Azure storage table top-level node. If you expand that node, you should be able to identify the Azure table according to the **Table Name** that you entered earlier on the **Azure** tab of the **New Session** dialog. Under this node, you will find all the Azure table fields that you can add as columns to the **Analysis Grid** viewer.

TIP

For analysis purposes, you can quickly create a view **Filter** based on the values of any field with the use of the **Add '<fieldName>' to Filter** item in the context menu that displays when you right-click a field in the **Details Tool Window**, just as you can for any Live Trace Session that has completed. You can also do this for the column entities in the **Analysis Grid** viewer. Filtering can isolate specific data that you might be looking for, which could be particularly useful when an Azure Storage table contains a high volume of data.

More Information

To learn more about Message Analyzer session configuration, see Starting a Message Analyzer Session.
To learn more about **Session Filters**, see Applying a Session Filter to a Data Retrieval Session.
To learn more about creating filter expressions, see Writing Filter Expressions.
To learn more about the **Field Chooser**, see the Field Chooser Tool Window topic.

See Also

Retrieving Azure Storage Blob Data

Loading System Event Log Data

2 minutes to read

Message Analyzer enables you to load system event data that is typically displayed in the **Microsoft Event Viewer**. The types of logs for which you can load data into Message Analyzer consist of **Applications and Services**, **Windows**, and others. To load data from a particular log, you simply select the log name on the **Event Logs** tab of the **New Session** dialog in the **Sources** list and then click the **Start** button to load the data into Message Analyzer. The interface from which you will work to load Event Log data into Message Analyzer is shown in the figure that follows.

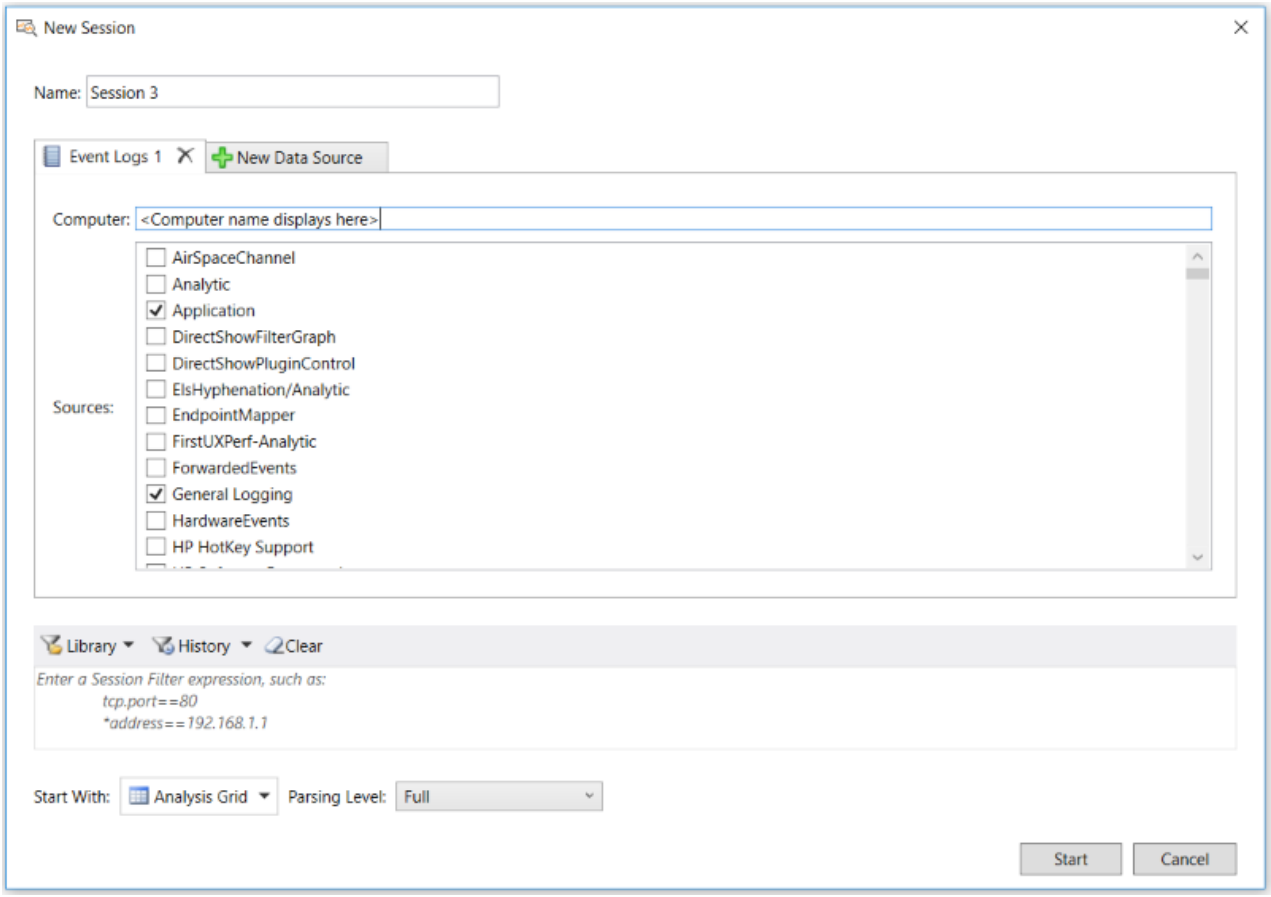

Figure 34: Event Logs data retrieval interface

After you load the data into the **Analysis Grid** viewer, you will typically see a row of data for each log entry, where the details of the log entry are contained in the **Summary** column of the **Analysis Grid**. If you select a row of data, you can view field names and values in the **Details Tool Window** below the **Analysis Grid** that correspond to the **Summary** column data. In addition, any diagnostic message that is associated with the selected entry is displayed in the **DiagnosisTypes** column of the **Analysis Grid** viewer for further examination.

To load data from a selected Event Log, perform the following procedure:

> **IMPORTANT**
>
> Before you perform the following steps, ensure that the **Event Viewer Import** preview feature is selected on the **Features** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. If not, select it and then restart Message Analyzer to enable the **Event Logs** option to appear in the **New Session** dialog under **Add Data Source**.

To load Event Log data into Message Analyzer

1. From the Message Analyzer **Start Page**, click the **New Session** button to display the **New Session** dialog.
2. Under **Add Data Source**, click the **Event Logs** button to display the **Event Logs** tab in the **New Session** dialog.
3. In the **Computer** text box, specify the name of the computer on which you wish to view **Event Logs**, which can be a remote or local computer. By default, the name of the local computer displays in this text box. Specify the computer name in the following format: "*computerName*", without the quotes and without spaces.
4. In the **Sources** list, place a check mark in the check box to the left of the Event Log name for which you want to view data.
5. When you finish with the input configuration, click the **Start** button in the **New Session** dialog to begin loading data from the selected Event Log into Message Analyzer.

> **TIP**
>
> You can also load system Event Log data into Message Analyzer from a *.evtx or *.xml file, if you save the former or export the latter from the **Microsoft Event Viewer**. In this case, you can load the data in these files through the **Add Files** feature in a Data Retrieval Session and examine the results in a chosen data viewer, which is typically the **Analysis Grid** viewer.

Deriving Input Data with PowerShell Scripts

5 minutes to read

This section describes two different methods that you can use to acquire data for input to Message Analyzer through PowerShell. The first shows you how to import data with a saved PowerShell script file that you target as a supported input file type (*.ps1) through a Data Retrieval Session. The script contained in such a file may invoke specific processes or functions that return data which you can view in Message Analyzer. The second method is similar, but instead enables you to use a PowerShell interface that is built into Message Analyzer to create a PowerShell query that returns its data to a viewer such as the **Analysis Grid**.

Importing Data Through a PowerShell Script File

As specified in Locating Supported Input Data File Types, Message Analyzer enables you to target PowerShell (.ps1) scripts as an input file type for a Data Retrieval Session. After you target one or more .ps1 files by adding them to the input files list on the **Files** tab of the **New Session** dialog, you can execute the scripts by clicking the **Start** button in the **New Session** dialog. The scripts might invoke other systems or functions to generate data, which will then be loaded through a Data Retrieval Session into Message Analyzer. Such PowerShell scripts can run on the local host or on remote target computers, as follows:

- **Local hosts** — you might have a PowerShell script that runs a cmdlet that returns data from the local Event Viewer, for example, `Get-EventLog -LogName Application -Newest 10`. This script gets the latest 10 messages in the Application log. If you use the **Add Files** feature of the **New Session** dialog to import a .ps1 file into Message Analyzer that contains this cmdlet, the data that it generates will be loaded into a Message Analyzer viewer that you specify (from the **Start With** drop-down list in the **New Session** dialog) after you start your Data Retrieval Session. Note that you also have the option to use the **Open** feature in the Message Analyzer **File** menu to quickly import a .ps1 file, execute its code, and load the data it generates into the default **Analysis Grid** viewer.
- **Remote hosts** — you can run a .ps1 script on a target remote host and return data to Message Analyzer in a manner that is similar to the local host scenario; however, some configuration is required, as described in the first two of the following steps:
  - Create a configuration file `psconfig.config` in the same directory where the remote .ps1 script exists.

    This file will contain the name of a single remote host from which to return data, using the format: `host=<hostname>`. Optionally, you can specify credentials in the following format: `user=<domain\username>` and `password=<Password>`. Specify each entry on separate lines in the .config file *without the angle brackets*. The default authentication method is Windows Integrated authentication.
  - Enable PowerShell remoting on the specified remote host by running the Enable-PSRemoting cmdlet.
  - Open the .ps1 file by using either of the methods specified in the "Local hosts" bullet point above.

Importing Data Through a PowerShell Query

To generate a PowerShell query that captures data from a particular local or remote host, you can construct the query with any cmdlets that you wish and display the output in Message Analyzer. This provides a convenient way of using PowerShell in the Message Analyzer environment where you can take advantage of Message Analyzer

analysis capabilities to review results. For example, you can use the **Analysis Grid** viewer to display the query output data in the **Summary** column of the grid. By selecting a message row in the grid, you can examine fields and values in the **Details Tool Window** that correspond to **Summary** column data. The figure that follows shows the data retrieval interface configuration with which you will work to load the data into Message Analyzer from the output of a PowerShell query.

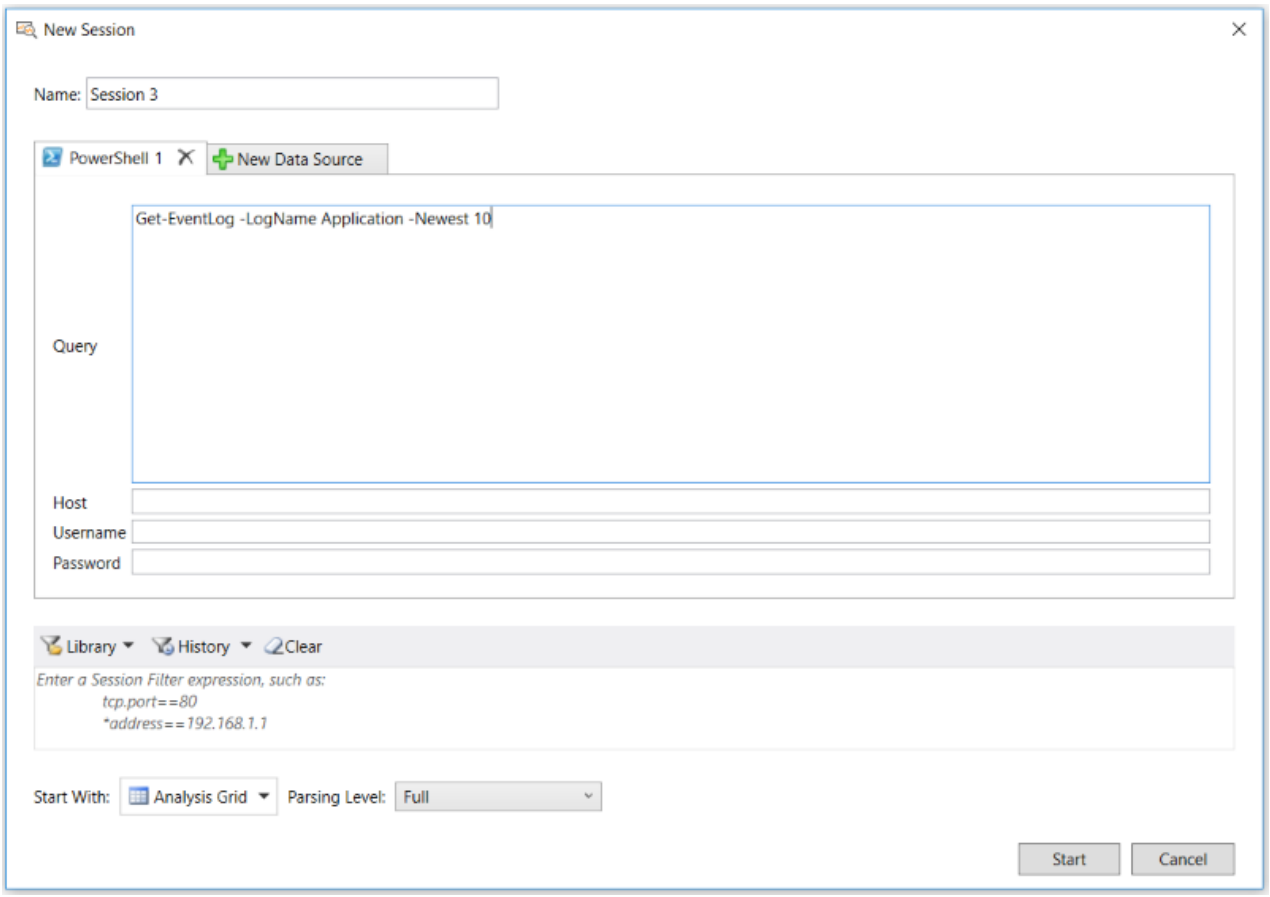

Figure 35: PowerShell Query data retrieval interface

To acquire data with a PowerShell query, perform the following steps:

> **IMPORTANT**
> Before you perform the following steps, ensure that the **PowerShell Import** preview feature is selected on the **Features** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. If not, select it and then restart Message Analyzer to enable the **PowerShell** option to appear in the **New Session** dialog under **Add Data Source**.

To acquire input data from a Power Shell query

1. On the Message Analyzer **Start Page**, click the **New Session** button to display the **New Session** dialog with the **Data Source** selection buttons.
2. Under **Add Data Source** in the dialog, click the **PowerShell** button to display the **PowerShell** query tab, from where you can specify the following:
   - **Query** — in this field, specify the PowerShell query code that targets the data you want to acquire. The figure above shows the hypothetical query `Get-EventLog -LogName Application -Newest 10`, which returns the latest 10 messages in the Application log of the local computer.
   - **Host** — to obtain target data from a remote host, you must specify the host name in this field in the following format: "*serverName*", without the quotes. If your query is for the local computer only, you can leave this field blank, as the PowerShell query that you specify will acquire data on the local host by default.

- **Username** — in this field, specify the user name that has appropriate permissions on the remote host where you are acquiring data, while using the following format: "*DomainName**Username*", without the quotes. If your query is for the local computer only, you can leave this field blank.
- **Password** — in this field, specify the appropriate password for the **Username** that you provided. If your query is for the local computer only, you can leave this field blank.

3. If not already selected, choose the **Analysis Grid** viewer from the **Start With** drop-down list in the **New Session** dialog as the viewer in which to assess your query results.

When you finish with the PowerShell query and host configuration, click the **Start** button in the **New Session** dialog to begin acquiring your data.

Loading SQL Data

2 minutes to read

Message Analyzer enables you to retrieve data from a SQL database table by using a built-in interface to provide connection information, SQL query code, and a reference timestamp. After you complete the input configuration and **Start** a Data Retrieval Session, you can display the data in the default **Analysis Grid** viewer in rows of information, with the SQL table data in the **Summary** column of the grid. The data in this column will be in a format that is typical of comma-separated value (CSV) or tab-separated value (TSV) delimiting. This means that you can view the values of the delimited data as separate fields in the **Details Tool Window** for any row of information that you select in the **Analysis Grid** viewer.

The figure that follows shows the interface from which you will work when loading SQL table data into Message Analyzer.

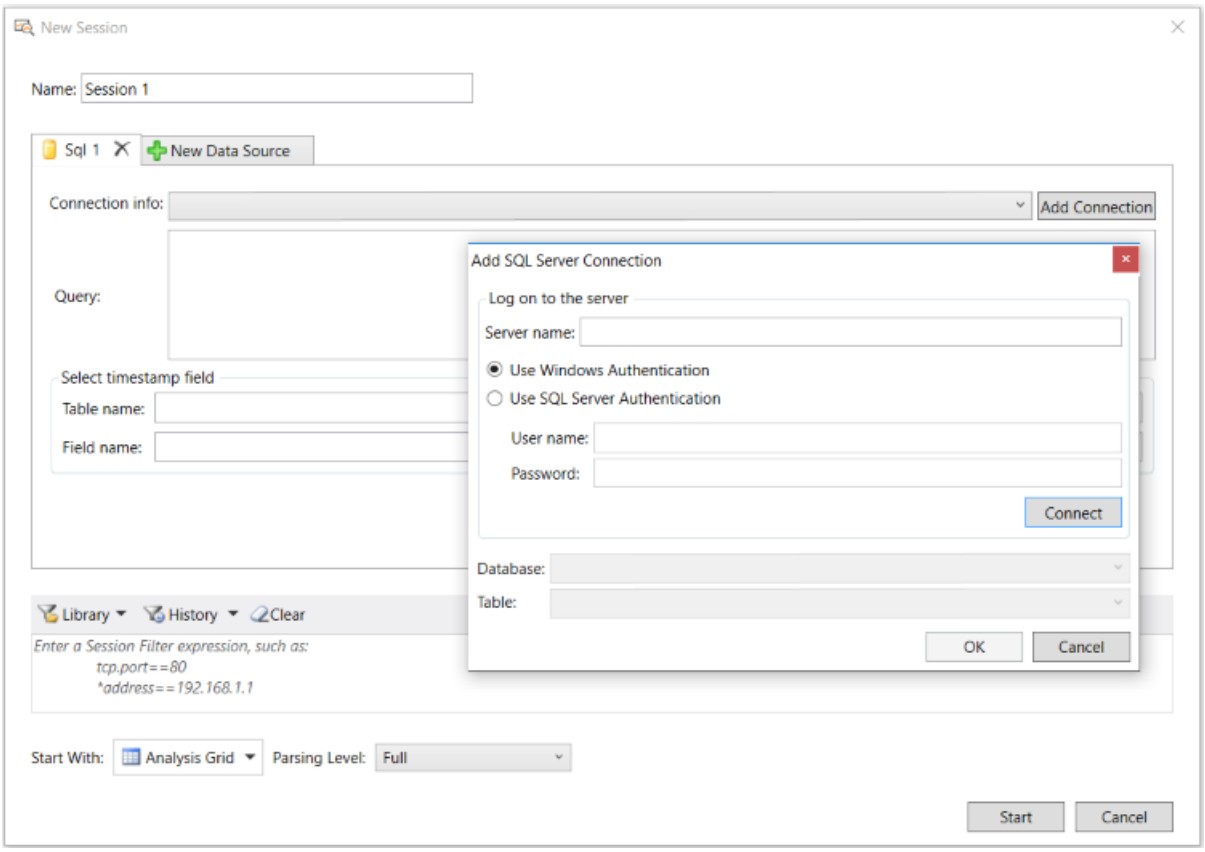

Figure 36: SQL table data retrieval interface

To load data from a specified SQL table, perform the following procedure:

> **IMPORTANT**
>
> Before you perform the following steps, ensure that the **SQL Table Import** preview feature is selected on the **Features** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. If not, select it and then restart Message Analyzer to enable the **Sql** option to appear in the **New Session** dialog under **Add Data Source**.

To load data from a SQL table into Message Analyzer

1. From the Message Analyzer **Start Page**, click the **New Session** button to display the **New Session** dialog.
2. Under **Add Data Source**, click the **Sql** button to display the **Sql** tab in the **New Session** dialog.

3. Create a SQL connection string by clicking the **Add Connection** button to display the **Add SQL Server Connection** dialog that is shown in the above figure. In the dialog, specify the following information:

   - **Server name** — specify the name of the SQL server you want to connect to by using the following format: "*serverName*", without the quotes.
   - **Authentication** — authentication options consist of the following:
     - **Use Windows Authentication** — the default selection, which uses your current logon credentials to authenticate to the SQL server.
     - **Use SQL Server Authentication** — specify this option if you want to provide an account with appropriate permissions on the server instead.

       If you use this option, then you will need to specify the **User name** in the following format: "*domain**Username*" without the quotes, along with the account **Password**.

4. Click the **Connect** button to validate a successful connection with your credentials, in which case, the **Database** and **Table** drop-down lists in the **Add SQL Server Connection** dialog will be populated with data.
5. From these lists, select an appropriate **Database** and **Table** from which you want to retrieve data.
6. When complete, click **OK** to create a new connection string, which should then appear and persist in the **Connection Info** drop-down list on the **Sql** tab of the **New Session** dialog.
7. In the **Query** box on the **Sql** tab of the **New Session** dialog, specify the SQL code that locates and manipulates the data you wish to extract from the previously specified **Database** and **Table**.
8. In the **Select Timestamp Field** pane on the **Sql** tab, select the appropriate **Table Name** and then specify the **Field Name** that contains the **Timestamp** you wish to use as a reference for your data.
9. When you finish with the input configuration, click the **Start** button in the **New Session** dialog to begin loading data from the selected SQL database table into Message Analyzer.

Loading WPP-Generated Events

11 minutes to read

Message Analyzer supports parsing and display of Windows software trace preprocessor (WPP)-generated events. Because these events make use of the ETW framework, Message Analyzer can capture them live or load them from a saved event trace log (ETL) file that is created by a system tool such as Logman or Netsh.

WPP is typically used by developers who want to generate events in their code for troubleshooting purposes. To do this, a developer embeds a simple function in code that, when executed, generates an ETW event. Unlike an ETW provider that relies upon a manifest to generate structured events for an ETW consumer, events generated by WPP are not defined by a manifest, but rather by a format contained in a PDB file, TMF file, configuration file, or a symbol server. Message Analyzer currently supports the first three of these.

Parsing WPP-Generated Events

To enable parsing of WPP-generated events, users must provide one of the following files to define the WPP event structure:

- **Program data base (PDB) file** — a symbol file for a trace provider that contains instructions for parsing and formatting the provider events. Unless you use the **Options** dialog method described in Specifying the Event Structure File Location to specify the path to your PDB file, then you must manually place the PDB file in the same directory as the ETL file, as described in Manually Locating a Directory. In addition, the PDB file must also have a name that is identical to the ETL file name, so that Message Analyzer can find the event structure information for parsing.
- **Trace message format (TMF) file** — a structured text file that contains instructions for parsing and formatting the events generated by a trace provider. Unless you use an XML configuration file or the **Options** dialog method described in Specifying the Event Structure File Location to specify the path to your TMF file, then you must manually place the TMF file in the same directory as the ETL file, as described in Manually Locating a Directory. In addition, the TMF file must also have a name that is identical to the ETL file name so that Message Analyzer can find the event structure information for parsing.
- **Configuration (.config) file** — an XML configuration file that enables you to create a store of one or more TMF files that parse events generated by different trace providers. The XML wpp.config file provides a <share|drive:> tag where you enter a share or drive location that contains one or more TMF files with event structure information, as indicated in Specifying the Event Structure File Location. Note that by using a configuration file, you can avoid the task of renaming TMF files, as described in the previous bullet point.

> **NOTE**
>
> If you only have a PDB file, you can use features on the **WPP** tab of the **Options** dialog to automatically convert it to TMF. If you want to use the XML configuration file to point to a share containing event structure information, you can use the command-line tool *Tracepdb* to manually convert the PDB file to a TMF file for use with the configuration file. Tracepdb.exe creates a trace message format (.tmf) file by extracting event formatting instructions from a PDB symbol file for a trace provider that used WPP software tracing macros. See Creating an XML Configuration File for further details.

The section that follows describes the methods you can use to specify the location of PDB and TMF files that contain event structure information for parsing a WPP-generated ETL file.

Specifying the Event Structure File Location

There are three different ways that you can specify the location of PDB and TMF files, the first of which is the easiest and overrides the others:

- **Use the Message Analyzer Options dialog** — specify the path to PDB and TMF files on the **WPP** tab of the **Options** dialog, as described in Using the Options Dialog.
- **Manually locate a directory** — place a PDB or TMF event structure file in a specific directory location, as described in Manually Locating a Directory.
- **Create an XML configuration file** — specify the share path to one or more TMF files in an XML configuration file that you create with specified XML code, as described in Creating an XML Configuration File.

Using the Options Dialog

The **Options** dialog now provides a dedicated UI feature on the **WPP** tab that manages WPP tracing symbol resolution. You can specify the path to a TMF file in the dialog and Message Analyzer will automatically reference such a file for the event structure required to parse a WPP-generated ETL file. In addition, if you only have a PDB file, you can specify the path to that file and the path to the command-line tool *Tracepdb.exe* on the **WPP** tab of the **Options** dialog and Message Analyzer will convert the PDB file to TMF for you.

You can display the **Options** dialog by clicking the global Message Analyzer **Tools** menu and the selecting the **Options** item. After displaying the **Options** dialog, click the **WPP** tab to use the features shown in the figure that follows:

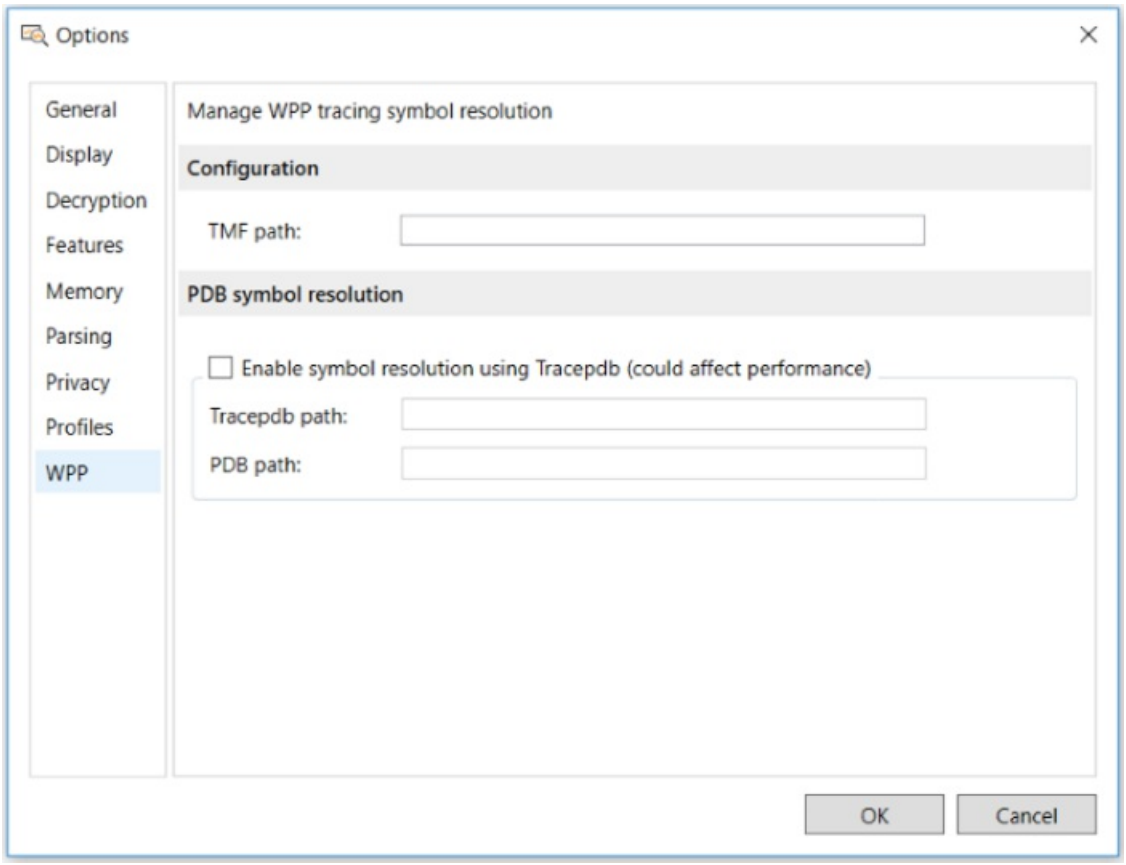

Figure 37: WPP tracing symbol resolution configuration

The list that follows describes the indicated features.

- **Configuration** section
  - **TMF path** — in this text box, specify the path to one or more existing TMF files, with each entry separated by a colon. Message Analyzer will create an OPN description from the information in each

file that you specify so that the WPP-generated ETL file/s can be parsed.

- **PDB symbol resolution** section

  You must select the **Enable symbol resolution using Tracepdb** check box to enable the following text boxes for data entry. Use this option if you have only a PDB file generated from the trace provider compilation or build process, in which case, Message Analyzer will convert the PDB file to TMF and create an OPN description for parsing WPP-generated ETLs.

  - **Tracepdb path** — in this text box, specify the path to the command-line tool *Tracepdb* to convert an existing PDB file to TMF. If you have a 64-bit installation of Windows Driver Kit 8.1 or Visual Studio, the path to Tracepbd is as follows: `C:\Program Files (x86)\Windows Kits\8.1\bin\x64\tracepdb.exe`

  More Information
  To learn more about *Tracepdb*, the commands that it provides, and download information, see Tracepdb.

  - **PDB path** — in this text box, specify the path to the PDB file that contains the event structure you want to convert to TMF for parsing a WPP-generated ETL file. Note that you can add multiple PDB paths in this text box if you have multiple ETL files to parse, but you must separate each path with a colon.

  For each PDB file that you specify, Message Analyzer generates and compiles an OPN description for parsing associated WPP-generated ETL files when loaded through a Data Retrieval Session.

When you click the **OK** button to exit the **Options** dialog, the WPP configuration immediately takes effect; however, OPN descriptions are created dynamically as described in Generating an OPN Description for PDB and TMF Files.

Manually Locating a Directory

To manually locate a directory for a PDB or TMF file and to make the file discoverable, you will need to do the following:

- Place the PDB or TMF file in the same directory as the target ETL file.
- Name the PDB or TMF file to match the name of the target ETL file.

  For example, if the name of your WPP-generated ETL file is events.etl, then your event definition file must be named events.pdb or events.tmf, as appropriate.

Message Analyzer will then be able to locate the appropriate directory location, given that it will know the directory from where you open the ETL file; it will also find the appropriate event definition file by matching the GUID in the ETL file with that of the PDB or TMF file that you are using.

Creating an XML Configuration File

You have the option to create a wpp.config file that specifies the path to one or more TMF files. You might want to use such a configuration file if you need to store multiple files that contain event structure information for parsing the events that were written by several different trace providers, or you might have multiple versions of the same TMF file for test purposes. The configuration file must contain the following XML:

```
<?xml version="1.0" encoding="utf-8"?>
<root>
<tmf>
<storeLocation>
<share|drive:>\events\wpp\WppTest\tmffile1
<share|drive:>\events\wpp\WppTest\tmffile2
  . . .
</storeLocation>

<versionMapping>
<startWith>6.0</startWith>
<folder>Vista</folder>
</versionMapping>

<versionMapping>
<startWith>6.1</startWith>
<folder>Windows7</folder>
</versionMapping>

<versionMapping>
<startWith>6.2</startWith>
<folder>Windows8</folder>
</versionMapping>
</tmf>
</root>
```

When using the wpp.config file, you must specify the directory location for all TMF files under the <store location/> tag in the XML above by using the <share|drive:> tag, rather than specifying PDB files, which are not supported in this configuration. If you only have a PDB file, you can use the command-line tool *Tracepdb* to convert the PDB file to a TMF file, as previously described. When this is the case, you should specify the TMFDirectory output path from the *Tracepdb* tool in the <share | drive:> tag of the XML code above.

Note that the <store location/> tag allows you to specify multiple TMF file versions, as shown in the code example above. Thereafter, you must place the wpp.config file in the following directory location. Note that you will need to create the "WppConfiguration" directory:

`%localappdata%\Microsoft\MessageAnalyzer\OPNAndConfiguration\WppConfiguration\`

> **NOTE**
>
> When you place your wpp.config file in this location, Message Analyzer will automatically search the store location to match the GUID of the appropriate TMF file with that of the ETL file you are attempting to parse.

Obtaining PDB and TMF Files

Depending on the tools you are using to log and display formatted trace events, you might be using either a TMF or PDB symbol file to store formatting information. Some tools require one or the other, while others can use either to extract the required formatting information.

A PDB symbol file for a WPP trace provider contains instructions for parsing and formatting events that are generated by the provider. Because these instructions are part of the trace provider source code, the WPP preprocessor can extract them from the code and add them to a PDB symbol file during compilation of a debug version of the trace provider or during the provider build process.

A TMF file for a WPP trace provider is a structured text file that also contains instructions for parsing and formatting events that are generated by the provider. A TMF file can be automatically generated by the trace provider build process. However, if for some reason you do not have a TMF file for your trace provider, you can generate one with the tool *Tracepdb*, as previously indicated, which extracts the formatting instructions from a PDB symbol file and then creates a TMF file to store them. You can perform this task manually at the command

line, or you can direct Message Analyzer to do it by making use of the configuration in the **WPP** tab of the **Options** dialog, as described earlier.

Generating an OPN Description for PDB and TMF Files

The WPP feature is enabled by default in Message Analyzer. When some component or trace provider code generates one or more ETW events from WPP, the events can be written to an .etl file by a system tool such as Logman or Windows Performance Monitor. The format that defines how to parse the events in the .etl file is stored in the previously mentioned PDB or TMF files. Message Analyzer uses an ETW adapter extension known as the WPP Import Adapter to process the events in an .etl file when you load one through a Data Retrieval Session or when you capture events in a Live Trace Session. However, before Message Analyzer can actually parse the event data, the WPP Import Adapter must generate an OPN description.

To enable Message Analyzer to parse WPP events, an OPN module needs to be generated so that the OPN description required to parse the WPP events can be compiled to the protocol object model (POM) in the PEF architecture and made accessible to PEF Runtime processing. The WPP Import Adapter dynamically generates the OPN module for the WPP trace provider and dynamically inserts WPP entry parameters as annotations in the individual WPP-generated events by locating such parameters in the PDB or TMF files.

Capturing WPP-Generated Events Live with Message Analyzer

For developers who want to use Message Analyzer to capture WPP-generated events live, you can do so if you have a corresponding managed object format (MOF) provider that is registered on your system. If it is registered, then Message Analyzer should expose the WPP/MOF provider in the **Available System Providers** list in the **Add System Providers** dialog. This dialog displays when you click **Add System Providers** in the **Add Providers** drop-down list on the toolbar of the **New Session** dialog during Live Trace Session configuration. You can then select the provider and run a trace to capture the events that are generated by the WPP/MOF-based trace provider. If the MOF provider is not registered on your system and you have the MOF file, you can manually register it by using the MOF compiler mofcomp.

More Information
To learn more about WPP, see WPP Software Tracing.
To learn more about PDB files, see PDB Symbol Files.
To learn more about TMF files, see Trace Message Format Files.

Loading OMS Log Data

4 minutes to read

Message Analyzer now enables you to load data from Operations Management Suite (OMS) logs so that you can take advantage of Message Analyzer data viewers and analysis capabilities. Message Analyzer provides a search interface to OMS Log Analytics that you can access with the **OMS** data source feature of the **New Session** dialog during Data Retrieval Session configuration. However, **OMS** is a preview feature and for it to appear as an input data source in the **New Session** dialog, you will need to select the **OMS Import** feature on the **Features** tab of the **Options** dialog and then restart Message Analyzer. The **Options** dialog is accessible from the global Message Analyzer **Tools** menu.

OMS Background Information

The data and resources that are managed by OMS are organized by an OMS Workspace, which you can think of as a unique OMS environment with its own data repositories and data sources. To access OMS data, you must have an Azure Subscription and credentials to log in to Azure. You can have multiple Workspaces in your Subscription to support multiple environments, for example, production and test. In a Workspace, you can enable different solutions that return operational data, for example, the Wire Data solution, which returns Network and performance data. You can now use Message Analyzer to extract OMS data from logs that are written by solutions that are enabled in a Workspace that you create. Message Analyzer does this by creating an interface for the search component of OMS Log Analytics, which enables you to write Azure Resource Manager (ARM) queries that access specific data for which you are searching in a particular OMS data collection in your Workspace.

For a single query, a user might access several log entries of different types in a single workspace, while the number of log entries actually returned by OMS depends on the query itself. For example, if you write an ARM query such as "events | top 100", you will return the first 100 events. If you specify a query such as "error", you will return all the log entries in your Workspace that contain errors.

More Information
To learn more about Log Analytics search syntax, see the Log Analytics Documentation topic on TechNet.

OMS Import Feature Process Flow

The processes that must occur in sequence to enable you to access data from OMS and display it in Message Analyzer consist of the following:

- A user **Login** provides the credentials that initiate the process of logging into Azure.
- The **Oms** tab of the **New Session** dialog displays a **Subscription** and one or more **Workspaces** in corresponding drop-down lists.
- The user selects his/her **Subscription** and a **Workspace** under that **Subscription**.
- The user specifies an ARM search query in the **Query** text box on the **Oms** tab.
- The user clicks **Start** to exit the **New Session** dialog, at which time Message Analyzer uses the token returned by the **Login** process to facilitate sending an ARM query to OMS.
- OMS returns one or more Java Script Object Notation (JSON) data structures to Message Analyzer for any log entry data that matches the input query.

- Message Analyzer parses the JSON and displays the data as individual messages, typically in the **Analysis Grid** viewer.

Using the Message Analyzer OMS User Interface

The interface for retrieving OMS log data is shown in the following figure.

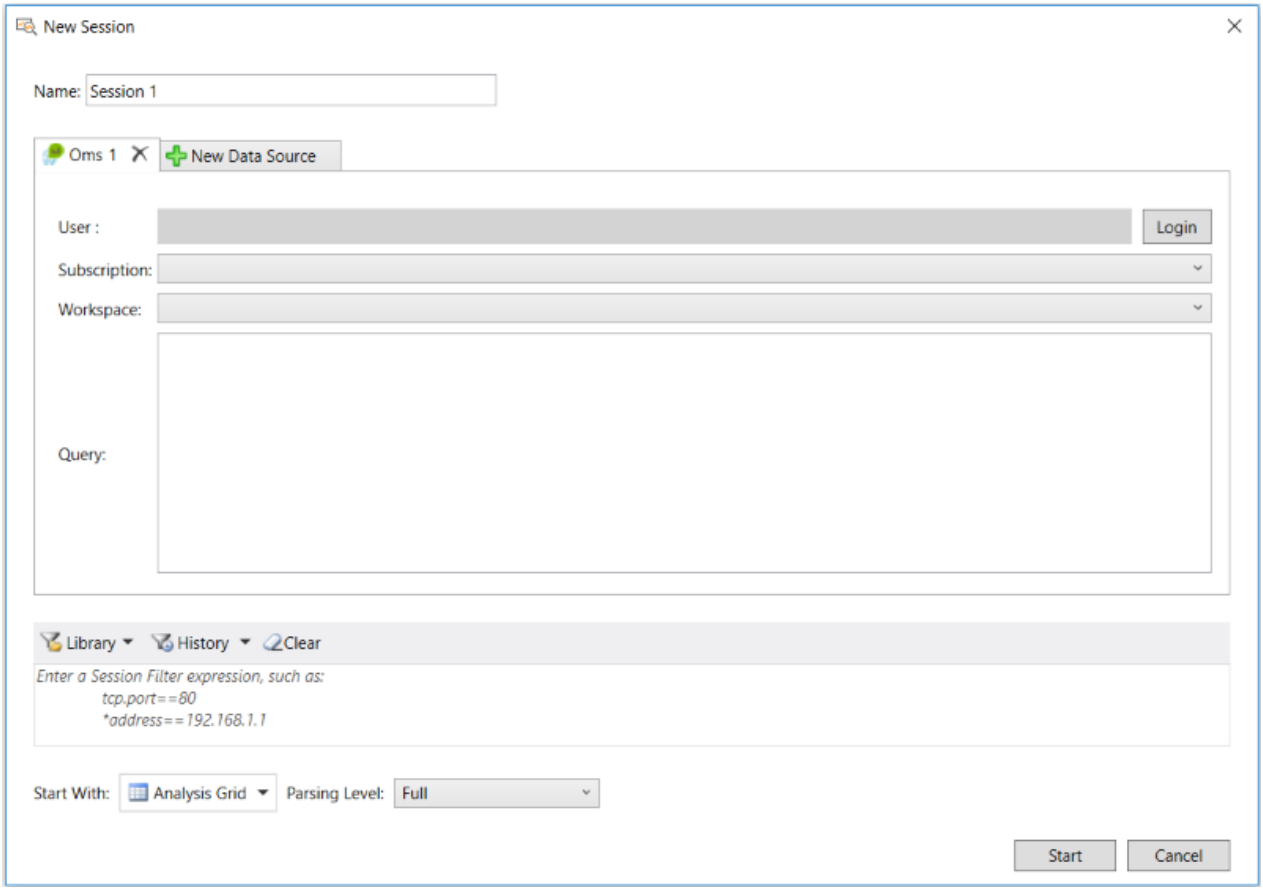

Figure 38: OMS log data retrieval interface

When you are ready to load data from OMS into Message Analyzer, perform the steps that follow while referring to the preceding figure for location of features.

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** application to launch Message Analyzer. To ensure that you have access to all features, run Message Analyzer as an Administrator by right-clicking the Message Analyzer application and selecting **Run as administrator** from the context menu that appears.

2. Ensure that the **OMS Import** feature is enabled in the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu.

   If the **OMS Import** feature is not enabled in the **Options** dialog, select the **OMS Import** check box and then restart Message Analyzer, as indicated earlier.

3. Click the **New Session** button on the global Message Analyzer toolbar to display the **New Session** dialog. The **Oms** button should appear under **Add Data Source** in the dialog.

4. Click the **Oms** button to display the **Oms** tab in the **New Session** dialog.

5. Click the **Login** button to specify your Azure account credentials and log in to Azure.

6. Click the **Subscription** drop-down and select your Azure Subscription in the list.

7. Click the **Workspace** drop-down and select an OMS Workspace in the list.

8. In the **Query** text box, write an ARM query for the data you want to retrieve.

9. In the **Start With** drop-down list, select the **Analysis Grid** viewer, if it is not already selected.
10. Click the **Start** button in the **New Session** dialog to begin data retrieval.
11. Observe that Message Analyzer displays the data requested by your query in the **Analysis Grid** viewer as individual messages.

More Information

To learn more about JSON, see the JSON Official Site.

Merging and Aggregating Message Data

3 minutes to read

A key feature of the Message Analyzer Browse-Select-View (BSV) model is that it enables you to merge the data from multiple files and aggregate data from multiple locations. This can be useful in scenarios where you are loading multiple log and/or trace files and you want to view and assess the data as a single message collection. If you start a Data Retrieval Session that loads data from multiple files at the same time, the data displayed by default in the **Analysis Grid** viewer will consist of an aggregation of chronologically sorted messages.

Merging Data From Multiple Sources and Locations

You can merge data from multiple sources by using the **Add Files** feature on the **Files** tab of a Data Retrieval Session. You can aggregate data from multiple locations by clicking the **New Data Source** tab in a Data Retrieval Session. Whenever you click this tab, a new **Files** tab displays, from where you can locate additional local or remote data sources. When you **Start** the Data Retrieval Session, the data from all sources is loaded into Message Analyzer and messages display in chronological order. You can also make use of the **New Data Source** tab when editing a session in the **Edit Session** dialog.

> **TIP**
> After you specify multiple data sources with the use of the **New Data Source** tab in Data Retrieval Session configuration and you load the data, you can add the **DataSource** field to the **Analysis Grid** viewer as a new column. You can do this by right-clicking the **DataSource** field under the **General** node of the **Field Chooser Tool Window** and then selecting the **Add As Column** command. You can then right-click the **DataSource** column in the **Analysis Grid** viewer and select the **Group** command to conveniently group the messages according to the data source name.
> In addition, if you have the **Grouping** viewer displayed, you can also add the **DataSource** field to this viewer as a new Group so you can organize the data by data source. You can do this by right-clicking the **DataSource** field in **Field Chooser** window and then selecting the **Add As Grouping** command. Thereafter, you can drag-and-drop the **DataSource** group label into the position occupied by the first group label in the current **Grouping** viewer **Layout**, as indicated by the red up and down arrows that display when a valid position is available on the **Layout** header. This action will organize your data into top-level **DataSource** groups, where the other groups of the displayed **Layout** will be sequentially nested subgroups.

Reconfiguring a Session to Aggregate More Data

To aggregate more data into an existing Data Retrieval Session for which results are being displayed, you can reconfigure it to add more saved data files. To begin, click the global Message Analyzer **Session** menu and then click the **Edit Session** command to display the **Edit Session** dialog. This dialog opens in the **Restricted Edit** mode as indicated in an information bar above the **Files** tab in the dialog. In this mode, you have only partial access to configuration features, which means that you can only add more files to the files list and **Apply** the change to load the messages contained in new file/s into the default data viewer. If you are loading a text log (.log file), you will also have access to the **Text Log Configuration** drop-down list for configuration file selection.

However, if you click the **Full Edit** button, you will have access to all the configuration features that are available for a Data Retrieval Session, such as a **Time Filter**, **Session Filter**, and check box selection/unselection of any files that existed in the session prior to reconfiguration. You should be aware that if you configure such features in **Full Edit** mode, it will trigger a reload of all the data from files that are selected in the files list when you **Apply** the changes. This is not critical, but performance could decline somewhat in this situation, depending on file sizes.

NOTE

If you initially loaded data from one or more target input files and you then reconfigure the session by specifying one or more additional input files in the **Edit Session** dialog, the message data from the additional files will be *appended* to the existing message data in the **Analysis Grid** viewer after you load the data. If you also change to the **Full Edit** mode in the **Edit Session** dialog and make other configuration changes, Message Analyzer will reload all the data and then display it in chronological order in the **Analysis Grid** viewer, providing that you are using this viewer.

See Also

Browse-Select-View Model
Editing Existing Sessions

Naming a Session

2 minutes to read

Although Message Analyzer provides a simple default name for each Data Retrieval Session, you can rename any session as a means of providing quick identification, for example, to reflect a particular issue on which you are working, the session configuration, or a data viewer that you used. You can specify a session name by entering one in the **Name** text box of the **New Session** dialog.

Identifying Sessions

Session names appear on all session tabs that display whenever you load data into Message Analyzer through a Data Retrieval Session or capture data through a Live Trace Session. Uniquely naming each session enables you to quickly identify and locate them in the **Session Explorer Tool Window** as named session nodes and as named session tabs below the global Message Analyzer toolbar during data analysis.

Changing Session Names

Prior to loading data into Message Analyzer, you can accept the default name that Message Analyzer provides for a Data Retrieval Session, or you can rename it at your discretion. After you load data into Message Analyzer, you have the option to change the name that you assign to a Data Retrieval Session by modifying it from the **Edit Session** dialog, which is accessible from the **Session** menu. Note that the session name can be different than the name of the file in which you save your session data. These behaviors also apply to a Live Trace Session.

Performing Data Retrieval

6 minutes to read

This section describes the methods that you can use to retrieve input data to load into Message Analyzer. This includes the quickest methods for doing so, such as the **Open** feature.

Retrieving Data Quickly

There are several methods that you can use to get data into Message Analyzer quickly. You can use these methods to find, load, and immediately display the contents of traces, logs, and other data file types that are described in Locating Supported Input Data File Types. These methods consist of the following:

- **Open** feature — the fastest way to retrieve saved data with Message Analyzer is to use the **Open** feature that is located on the global Message Analyzer toolbar as a folder icon that contains a drop-down list, or you can use the **Open** command that is accessible from the Message Analyzer **File** menu. From either of these locations, you can select the **From File Explorer** command from the **Open** drop-down list to display the **Open** dialog, from where you can navigate to and select the files that contain the data of interest.

> **NOTE**
> If you use the **Open** feature, you can load data from multiple files simultaneously.

> **TIP**
> You can also select the **From Other File Sources** command from the **Open** drop-down list, to display the **File Selector** dialog, from where you can target Azure logs as input to Message Analyzer.

- **Drag-and-drop** method — you can retrieve saved data quickly by dragging-and-dropping one or more saved trace or log files onto various Message Analyzer locations, for example, the **Files** tab of a Data Retrieval Session, the **Session Explorer Tool Window**, or the **Start Page**.

> **IMPORTANT**
> If you elect to run Message Analyzer in Administrator mode, it can result in varying security contexts between applications. This means that you will be unable to use the drag-and-drop feature to open saved trace and log files in this mode.

- **Windows Explorer** method — you can use **Windows Explorer** application to navigate to and select the saved files from which you want to retrieve data.
- **Recent Files** feature — the global Message Analyzer **File** menu contains a **Recent Files** list that enables you to quickly display data and resume your work from earlier sessions that you saved. You can quickly open one of the files in the list by clicking it and Message Analyzer will then display its data in the current default viewer.
- **Ctrl+O** keyboard shortcut — you can use the keyboard shortcut `Ctrl+O` to display the **Open** dialog.

When you use these features to load data into Message Analyzer, the data loading process is immediately invoked in the background and the results quickly begin to display in the default data viewer, for example, the

Analysis Grid. For this to occur, Message Analyzer automatically creates a new Data Retrieval Session that is implicitly named based on a default session name, which you can edit to customize as you wish. However, when using the previously described input methods, you can only edit the session name if you click the **Edit Session** icon on the global Message Analyzer toolbar *after* the data has loaded into Message Analyzer. The only exception to the automatic startup of a Data Retrieval Session using the previously described input methods is if you have selected a .log file for input to Message Analyzer. In this case, you will need to provide some additional configuration, as described in the next section.

Specifying a Text Log Configuration File

If you are loading data from a .log file, Message Analyzer opens the **New Session** dialog so that you can specify a built-in **Text Log Configuration** file to parse the log. However, if you have previously specified a particular configuration file as the default on the **General** tab of the **Options** dialog (accessible from the global Message Analyzer **Tools** menu), this does not occur. Rather, Message Analyzer immediately proceeds to load your log data without opening the **New Session** dialog. In addition, if you cannot find a built-in configuration file that will successfully parse your log file, you may need to create one, as described in Opening Text Log Files.

Loading Data with a Specified Session Configuration

You can also navigate to trace files or logs that contain the data you want to work with in a Data Retrieval Session, by clicking the **Add Files** button on the **Files** tab of the **New Session** dialog to display the **Open** dialog. After you select files and click **Open** to exit this dialog, the files display in the files list on the **Files** tab of the **New Session** dialog. At this point, no data is loaded into Message Analyzer, as the files list is only a set of pointers to the files that contain the data you are targeting.

To continue with Data Retrieval Session configuration prior to loading data, you have the option to select specific files in the files list to create a unique collection of messages that you want to analyze, as described in Selecting Input Files for Data Retrieval. Prior to starting your Data Retrieval Session, you also have the option to select a different data viewer from the **Start With** drop-down list in the **New Session** dialog; otherwise, your data will be displayed in the data viewer that is set as the default in the **Session Viewer** section on the **General** tab of the global **Options** dialog that is accessible from the global Message Analyzer **Tools** menu. In addition, if you want to filter the data you are loading to retrieve specific information to focus on, you can do so by specifying a **Session Filter** or a **Time Filter**, as described in Selecting Data to Retrieve.

Selecting Input Files for Data Retrieval

When configuring a Data Retrieval Session, you can specify the files containing the target data that will display in a Message Analyzer viewer, by placing a check mark in the check box next to the file names in the files list on the **Files** tab of the **New Session** dialog. Message Analyzer keeps track of the **Message Count** for each file that you add, the number of files available for data retrieval, the number of files you actually selected, and the total number of aggregated messages from all the files selected in the list. Other file attributes are also displayed for each file, such as the **Name**, **Size**, **File Type**, **Start Time**, and **End Time**. The current exceptions to this are .etl and .pcap files, for which message count and time ranges may not be displayed in the Data Retrieval Session configuration.

Caution

Loading data from very large input files can impact the performance of the data loading process, especially on 32-bit computers.

NOTE

If Message Analyzer detects that an input file from which you are loading data was processed with an out-of-date parser, you will be prompted to reparse that data with an updated parser. If you choose not to reparse, Message Analyzer will not open the file. If you do reparse, note that any **Bookmarks** and **Comments** in the file will be preserved, but in some cases they might be relocated to the lowest stack message due to OPN revisions that might have impacted the origins tree (messages underlying top-level).

Starting a Configured Data Retrieval Session

When you are ready to load the data into Message Analyzer from a *configured* Data Retrieval Session and to display it in the default data viewer, click the **Start** button in the **New Session** dialog.

More Information

To learn more about using the **File Selector** and loading data from Azure logs, see Retrieving Azure Storage Blob Data.

Procedures: Using the Data Retrieval Features

32 minutes to read

The procedures in this section encapsulate some of the main functionalities described in Retrieving Message Data. They serve as simple examples that demonstrate how to use Message Analyzer features to retrieve saved logs and files in the most efficient manner. These procedures, rather than serving as troubleshooting scenarios, also demonstrate the use of some data analysis tools that manipulate loaded data.

> **NOTE**
>
> Although these procedures demonstrate the use of Message Analyzer capabilities in some basic scenarios, they are only a sampling of what you can accomplish with Message Analyzer, given that you can also apply the methodologies described here to many other scenarios.

Procedure Overviews

A brief description of each procedure is included here for review, as follows.

Load and Display Saved Data — shows how to browse for saved files that contain a message collection you want to load into Message Analyzer through a Data Retrieval Session and display it in a selected data viewer.

Select Specific Data from a Saved Trace File — shows how to select specific data from a saved file by applying a **Session Filter** to the data loading process via a Data Retrieval Session.

Display Different Data Viewers for Session Results — shows how to view results in selected data viewers that provide different presentation formats to enhance your data analysis perspectives.

Load Saved Data with the Open Feature — shows how to quickly load and display data from a saved file by using the **Open** feature.

Load Saved Data From Recent Files or Drag-and-Drop — shows how to quickly load and display data from a saved file by using the **Recent Files** list. Also describes the use of drag-and-drop as an alternate method for opening files.

Apply a Time Filter to Data Loading and Save the Message Collection — shows how to load a message collection from multiple input files with a **Time Filter** applied; and how to save it to a single file in the default Message Analyzer .matp format.

Load Saved Data from a Text Log — demonstrates how to load data from a textual .log file into Message Analyzer with the use of a built-in **Text Log Configuration** file.

Retrieve Data from Log Files in Azure Storage BLOB Containers or from an Azure Table — demonstrates how to access, load, and view log data that is stored in Azure Storage binary large object (BLOB) containers or data that exists in an Azure table.

IMPORTANT

At least one procedure in this section makes use of the drag-and-drop feature. If you have not logged off Windows after the first installation of Message Analyzer, please log off and then log back on if you wish to use drag-and-drop. This action ensures that the subsequent logon that follows installation has the appropriate privileges from the Message Capture User Group, which in turn enables Message Analyzer to have the same security context as Windows Explorer; otherwise, drag-and-drop will not work.

Also note that if you run Message Analyzer as an Administrator, the security context between applications that is established prevents drag-and-drop from working properly.

Load and Display Saved Data

In the following procedure, you will load saved trace data through a Data Retrieval Session and display it in the Message Analyzer default **Analysis Grid** viewer.

To use a Data Retrieval Session to load and display saved data

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.
2. Click the global Message Analyzer **File** menu, point to **New Session**, and then click **Files** in the **New Session** submenu to display the **New Session** dialog for Data Retrieval Session configuration.
3. On the **Files** tab toolbar of the **New Session** dialog, click **Add Files** to launch the **Open** dialog and then navigate to the file/s that contain the data of interest.
4. In the **Open** dialog, select the file/s containing the data you want to load and then click **Open** to exit the dialog.

   TIP

   Comparing data from a Live Trace Session with associated data that is loaded from a Data Retrieval Session provides a convenient method for analyzing current and historical data.

5. In the files list that displays, ensure that there is a check mark in the check box next to the file/s that contain the data you want to load, or alternatively, remove the check mark from files that contain data which, for the moment, you do not want to load.

   If you have a large number of files in the list and you need to search for suitably named files that contain specific data, as described in Naming Saved Files, enter the appropriate file name characters in the search box on the **Files** tab toolbar to highlight them in the list.

6. Ensure that the **Start With** drop-down list in the **New Session** dialog specifies the default **Analysis Grid** viewer, or alternatively, select a different viewer.
7. Click the **Start** button in the **New Session** dialog to begin loading data from the selected files into Message Analyzer.

   The loaded data displays in the specified Message Analyzer data viewer.

8. To load data from one or more additional files, for example, files that you *unselected* in the initial data loading configuration, click the **Edit Session** icon on the global Message Analyzer toolbar to open the **Edit Session** dialog and expose the configuration of the current Data Retrieval Session, as you originally specified it.
9. Select additional files in the files list or click **Add Files** to locate and add one or more files to the files list, and then place a check mark in the check box next to each file that contains the data you want to load into

your existing message collection. By default, the **Restricted Edit** mode in which the **Edit Session** dialog opens only enables you to add more files to the files list, since other configuration features are disabled in this mode.

> **TIP**
> If you add a check mark to the **Select Added Files** check box on the toolbar of the **Files** tab, all files that you add to the files list with the **Add Files** feature are automatically selected for inclusion in the data loading process.

10. Click the **Apply** button in the **Edit Session** dialog to load the new data into the previously specified Message Analyzer data viewer.

    If your data viewer is the **Analysis Grid**, note that the new data you are adding is appended to the existing set of messages in the tree grid of this viewer.

> **TIP**
> When analyzing data that you loaded from multiple input data sources, as described in Configuring Session Scenarios with Selected Data Sources, you have the option to organize and summarize the loaded data into Groups that are labeled by data source name. To do this, locate the **DataSource** field in the **General** category of the **Field Chooser Tool Window** and then execute the **Add as Grouping** command by selecting it from the context menu that displays after you right-click the **DataSource** field.

Select Specific Data from a Saved Trace File

In the following procedure, you will select specific trace data to load into Message Analyzer through a Data Retrieval Session, by applying a **Session Filter** to the data loading process.

To select specific data in a Data Retrieval Session

1. Perform steps 1 and 2 of the procedure in Load and Display Saved Data.
2. On the **Files** tab toolbar, ensure that there is a check mark in the **Select Added Files** check box, so that all files that you add to the files list with the **Add Files** feature are automatically included in your data loading configuration.
3. On the **Files** tab toolbar, click **Add Files** to launch the **Open** dialog and then navigate to the file/s that contain the trace data of interest.
4. In the **Open** dialog, select the file/s containing the data you want to load into Message Analyzer and then click **Open** to exit the dialog.
5. From the **Library** drop-down list on the toolbar above the **Session Filter** text box, select a built-in Filter Expression, or alternatively, create your own custom Filter Expression to specify the filtering criteria that you want to apply to the input messages that are to be loaded into Message Analyzer.

   For example, you might specify a simple expression such as `IPv4.Address==192.168.1.1`, to filter for messages that contain a specified IP address only. You can also select a recently used filter from the **History** drop-down list on the previously indicated toolbar in the **New Session** dialog. In this example, the IP address in italics is a placeholder for an actual IP address that you will include in this filter.
6. Optionally, select a data viewer from the **Start With** drop-down list of the **New Session** dialog, or simply accept the default data viewer, for example, the **Analysis Grid**.

> **NOTE**
> At any time prior to loading data in the next step, you have the option to remove the check mark from any files that contain data you do not want to load at the moment.

7. Click the **Start** button in the **New Session** dialog to begin loading your selected data into Message Analyzer.

8. At your discretion, you can return to the configuration of your Data Retrieval Session to specify a different Filter Expression or message collection configuration for loading data into Message Analyzer, by clicking the **Edit Session** icon on the Message Analyzer global menu.

   The **Edit Session** dialog opens in the **Restricted Edit** mode, with an information bar that specifies the following:
   Add new files or data sources without causing a data reload. Other configuration changes in Full Edit mode cause a reload of all data.
   This means that you can add new data files into the target files list and load that data without Message Analyzer having to also reload the data from the existing files. But any other configuration changes that you specify for the current Data Retrieval Session will cause Message Analyzer to reload all data, which includes messages from all files that represent the original message collection, in addition to messages from any new files you are adding. When this occurs, you might notice slower performance as Message Analyzer reloads the data.

9. In the **Edit Session** dialog, click the **Full Edit** button to enable all configuration features for your Data Retrieval Session.

10. In the **Edit Session** dialog, select a different **Session Filter** from the centralized Filter Expression **Library**, for example, `#DiagnosisTypes==2`. If you select this filter, you will load and view only the messages that contain validation errors, for analysis purposes. A validation error is an indication that a message does not align with its protocol definition, as described in the Diagnosis Category topic.

11. After you have modified the Data Retrieval Session by selecting a different **Session Filter** from the centralized **Library**, click the **Apply** button to view the results of the new Data Retrieval Session configuration in the data viewer that you initially selected.

    Note When modifying an existing Data Retrieval Session, you cannot change the data viewer in which to display your data, as this capability is disabled in the **Edit Session** dialog. As a result, the messages that load from your modified Data Retrieval Session after you click the **Apply** button will continue to display in the data viewer that you initially specified. If you want to see the loaded data in a different data viewer, you can select one from the **New Viewer** drop-down list on the global Message Analyzer toolbar, as described in the procedure that follows.

More Information
To learn more about the **Edit Session** dialog, see Editing Existing Sessions.

Display Different Data Viewers for Session Results

In the following procedure, you will display different views of a loaded message collection by launching other data viewers. Doing so can help you obtain different analytical perspectives when assessing your data.

To display data in different viewers

1. Start a Data Retrieval Session with the default **Analysis Grid** viewer and load data into Message Analyzer by following the steps of one of the previous procedures, as appropriate.

   Thereafter in this procedure, you will be selecting other viewers.

2. Observe that the loaded data displays in the Message Analyzer **Analysis Grid** viewer, or whichever viewer is set as the default.

3. Confirm that the **Session Explorer Tool Window** is open. If not, open it by selecting the **Session Explorer** item in the **Windows** submenu that is accessible from the global Message Analyzer **Tools** menu.

4. To create a different view of your data, right-click anywhere in the **Session Explorer** window and highlight **New Viewer** in the context menu, highlight **Charts (Deprecated)**, and then select the **Protocol Dashboard** item. Alternatively, you can select the **Protocol Dashboard** viewer from the same location in the **New Viewer** drop-down list that is accessible from the global Message Analyzer toolbar.

5. Observe that the loaded data displays in the **Protocol Dashboard** viewer as a separate viewer tab in the same session.

   The **Protocol Dashboard** viewer contains several graphic data visualizers that include **Top Level Protocols Summary** components in **Table** grid, **Bar** element, and **Pie** slice formats that correlate message count versus module/protocol types, in addition to a **Top Level Protocols Over Time** chart in X-Y axis format that displays message count per module/protocol versus the trace timeline.

6. To create a another view of your data, right-click anywhere in the **Session Explorer** window and highlight **New Viewer** in the context menu, and then select the **Top Talkers Top 20** view **Layout** from the **Chart** drop-down in the **New Viewer** drop-down list,.

   > **NOTE**
   > This **Layout** contains a **Bar** element chart that displays the distribution of traffic volume for each IPv4 or IPv6 conversation that occurred across the timeline of a set of trace results. It also provides a relative indication of the percent bandwidth consumed by each conversation with respect to the total message count in the trace.

7. Right-click again anywhere in the **Session Explorer** window, highlight **New Viewer**, and then select the **Pattern Match** item.

8. Execute a Pattern Expression in the **AVAILABLE PATTERNS** pane, by placing a check mark in a chosen Pattern Expression check box.

   For example, you could select the **TCP Three-Way Handshake** check box to view and analyze the messages — respectively sent from and received by source and destination nodes — that successfully participated in TCP connection handshake communications across the trace timeline.

   > **TIP**
   > Because the **Pattern Match** viewer can interact with the **Analysis Grid** viewer, it may be useful to redock these viewers so that you can more effectively see the results of **Pattern Match** viewer message selection as it displays in the **Analysis Grid** viewer. See Working with Message Analyzer Window Layouts for more information.

9. To create a another view of your data, right-click anywhere in the **Session Explorer** window and highlight **New Viewer** in the context menu, and then select the **Gantt** viewer.

   This viewer shows the distribution of messages in each IP conversation across the timeline of a trace for each protocol that generated such messages. The protocols are color-coded in a **LEGEND** for quick identification, where you can select the protocol messages you wish to view.

10. Repeat these steps to specify different data viewers as needed. For example, you might specify one of the many other **Layouts** for the **Chart** viewer, which exist in the **Message Analyzer Chart View Layouts** asset collection. Note that this collection is accessible from the **New Viewer** drop-down list, wherever it appears, in every Message Analyzer installation.

11. To enhance your analysis perspectives, poll through the various data viewers by clicking the session nodes for each viewer type in the **Session Explorer** window, for the current Data Retrieval Session. Note that you can also manually select the corresponding session tabs to look at the data format of each viewer.

More Information
To learn more about the **Protocol Dashboard** viewer, see the Protocol Dashboard topic.
To learn more about the many **Layouts** that you can display for the **Chart** viewer, see the Chart Viewer Layouts topic.
To learn more about Pattern Expressions, including the built-in Pattern Expressions, see the Pattern Match Viewer section.
To learn more about the **Gantt** viewer, see the Gantt Viewer topic.

Load Saved Data with the Open Feature

In the procedure that follows, you will display data quickly from a saved trace or log file by using the **Open** feature.

To quickly open a saved file and display its data

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.

2. Click the global Message Analyzer **File** menu, select **Open**, and then click the **From File Explorer** item to launch the **Open** dialog. You can also use the keyboard shortcut `Ctrl+O` to open the dialog.

3. Navigate to the saved file containing the data you want to display, select the file, and then click **Open**.

   The saved data displays in the default data viewer, for example, the **Analysis Grid** viewer.

Caution
If you load a custom *.log file through the **Open** feature, Message Analyzer will open the **New Session** dialog first to display the configuration for a Data Retrieval Session. This action enables you to select a **Text Log Configuration** file that will fully parse the log data after you **Start** the session, otherwise, it is likely that Message Analyzer will be unable to parse the data in your custom log data.

The exception to this is when you have a default configuration file already specified in the **Text Log Files** pane on the **General** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. When this is the case, Message Analyzer does not display the **New Session** dialog, but rather, automatically begins parsing and loading the data from the *.log file into the default data viewer. Note that Message Analyzer will *fully* parse the .log file data only if the right configuration file is specified, which might be either a custom configuration file that you created, or one of the default configuration files that is provided with every Message Analyzer installation. For more information, see Opening Text Log Files.

> **NOTE**
> If you load a trace file that was saved with one or more out-of-date parsers, Message Analyzer prompts you to reparse the trace.

Load Saved Data From Recent Files or Drag-and-Drop

In the procedure that follows, you will quickly load and display trace data by using the **Recent Files** feature or the drag-and-drop feature, as alternate methods for loading data into Message Analyzer. To ensure that drag-and-drop functions properly, you will need to have logged off and back on Windows at least once since you installed Message Analyzer, as previously indicated in this section. In addition, ensure that you do not run Message Analyzer as Administrator if you want to use the drag-and-drop feature. If you do, the inconsistent security

contexts of Message Analyzer and **Windows/File Explorer** will prevent drag-and-drop from working properly.

To quickly load and display saved data

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.
2. Click the global Message Analyzer **File** menu and then select **Recent Files** to display the list of recently accessed trace and log files.
3. In the **Recent Files** submenu, click a file that contains data that you want to quickly display in the default data viewer, providing that you have one or more files in the list.

   > **NOTE**
   > If you load a text .log file through the **Recent Files** feature, the same behavior that is described in the previous **Caution** applies.

4. Alternatively, drag-and-drop one or more saved trace files from **Windows/File Explorer** to any location on the Message Analyzer user interface, such as the **Start Page**, global toolbar, or onto the **Session Explorer Tool Window**.

   Message Analyzer immediately displays the trace file data in the default session viewer, providing that reparsing is not required, in which case you are advised that Message Analyzer needs to reparse the data. If you drag-and-drop more than one trace file, the data for each file is appended in a single session viewer tab.

   Note that you can also drag-and-drop saved .log files onto the files list on the **Files** tab of the **New Session** dialog. This can be useful if you know in advance that you will need to specify a **Text Log Configuration** file in the dialog.

Apply a Time Filter to Data Loading and Save the Message Collection

In the procedure that follows, you will load a message collection consisting of data from multiple files while applying a **Time Filter** and you will save the results to a single file in the default Message Analyzer .matp file format.

To apply a Time Filter to the data loading process and save the message collection

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.
2. Click the global Message Analyzer **File** menu, point to **New Session**, and then select **Files** in the **New Session** submenu to display the **New Session** dialog for a Data Retrieval Session.
3. On the toolbar of the **Files** tab, click **Add Files** to launch the **Open** dialog and then navigate to the file/s that contain the data you want to load into Message Analyzer.
4. In the **Open** dialog, select the file/s containing the data you want to load and then click **Open** to exit the dialog.

   This might consist of any combination of .matu, .matp, .cap, .log, .etl, or other files that contain message data that was captured or logged in a similar time frame. For example, this could be data from several large log files that you want to aggregate and load into Message Analyzer for analysis purposes.

5. In the files list that displays, ensure that you have a check mark in the check box next to the file/s that contain the data you want to load.
6. In the **Time Filter** pane on the **Files** tab of the **New Session** dialog, adjust the left and right time slider controls to define a window of time in which to view data.

For example, if you have one or more large input files with target data, you might want to focus on a particular time slot in which you suspect that a particular issue has occurred to minimize consumption of system resources, rather than load all the message data contained in the input files. You can do this with a **Time Filter**, which loads only the messages with timestamp values that fall within a specified window of time.

> **NOTE**
> If you have a collection of target input files, the **Start Time** and **End Time** values that display in the **Time Filter** configuration are inclusive of the earliest and latest chronological time value, respectively, that is detected in any input file in the files list.

7. Optionally, select or configure a **Session Filter** for your Data Retrieval Session to isolate specific information that you want to focus on, as follows.

   If you want to be even more specific about the data that you load from target input files into Message Analyzer, you can specify a Filter Expression in the **Session Filter** text box of the **New Session** dialog, either by selecting a built-in filter from the centralized Filter Expression **Library**, or by configuring one manually.

   For example, for an input trace file, you might add a filter such as: `IPv4.DestinationAddress == 192.168.1.1`, to load the traffic that went to or from the specified address only; or for a log file, you might use a filter such as: `*Summary contains <"searchString">`, to exclude all messages except those that contain a specified string in the **Summary** column of the **Analysis Grid** viewer.

8. Click the **Start** button in the **New Session** dialog when you are ready to load the data.

   The data from the target files that you specified in the Data Retrieval Session configuration is filtered and loaded into Message Analyzer; it then displays in the default data viewer, for example, the **Analysis Grid**.

9. Create different data analysis perspectives by applying various Message Analyzer data assessment and analysis tools, as follows.

   After the data is loaded, you can optionally apply additional data manipulation techniques to further analyze or isolate specific messages of interest for enhancement of your analytical perspectives. For example, you can do this with any of the following operations.

   - Sorting columns.
   - Adding new columns with the **Field Chooser** to display other field data of a protocol or module that exists in your trace results. See Field Chooser Tool Window for further details.
   - Executing a right-click **Group** command on one or more **Analysis Grid** viewer columns to group data, for example, the **DiagnosisTypes** column. See Using the Analysis Grid Group Feature for further details.
   - Specifying **Column Filters** in the column search boxes that appear when you click the **Show Column Filter Row** icon next to the **MessageNumber** column of the **Analysis Grid** viewer. See Filtering Column Data for further details.
   - Specifying a **Viewpoint** from the Filtering toolbar above the main analysis surface where viewer data displays. See Applying and Managing Viewpoints for further details.
   - Specifying a view **Filters** from the Filtering toolbar. See Applying and Managing Filters for further details.
   - Choosing different viewer **Layouts** that focus on the data of different fields, which includes **Chart**

viewer **Layouts** that contain different types of visualizer components, in **Bar** element, **Pie** slice, **Timeline**, or **Table** grid formats. See Applying and Managing Analysis Grid Viewer Layouts, Grouping Viewer Layouts, and Chart Viewer Layouts for further details.

- Selecting different data viewers that enhance your analytical perspectives. See Data Viewers for further details.
- Applying a removable **Time Filter** from the Message Analyzer Filtering toolbar to further define the window of time in which to view data. See Applying a Time Filter to Session Results for further details.

10. Save your data by using the **Save/Export Session** dialog, as follows.

    When you have a set of messages that exposes a particular issue you are working on, such as a group of errors that might have occurred in some module at a specific source or destination address; or if you simply need to save your data to resume analysis later on, you can do so by clicking the **Save** item in the global Message Analyzer **File** menu or by clicking the **Save** icon on the global Message Analyzer toolbar to display the **Save/Export Session** dialog, and then doing one of the following:

    - Select the **All Messages** option to save all the data in a message collection.
    - Select the **Filtered Messages** option to save a message collection to which a view **Filter** was applied.
    - Choose the **Selected Messages** option after selecting/highlighting one or more messages with your mouse.

    Note that you have the option to save a message collection in the Message Analyzer native .matp file format, or you can export to a .cap file for use in other applications. See Compatibility with Exported CAP Files for more information about .cap file interoperability with other network troubleshooting tools.

11. From the **Save As** dialog, navigate to the directory location where you want to save the selected message data.
12. In the **File name** text box of the **Save As** dialog, specify a name for the message file, or use the default name if one displays.
13. Click **Save** when finished.

More Information
To learn more about manipulating message data for analysis purposes, see the Analysis Grid Viewer, Common Data Viewer Features, and Procedures: Using the Data Filtering Features topics.
To learn more about how to create an input window of time configuration in which to view retrieved data, see Applying an Input Time Filter to a Data Retrieval Session.

Load Saved Data from a Text Log

This procedure shows you how to load data from a text-based .log file. To successfully parse a text log, Message Analyzer requires an OPN configuration file to parse the fields of data in the log. After you load a *.log file into the **New Session** dialog for a Data Retrieval Session, you will find a **Text Log Configuration** drop-down list that enables you to select from a collection of built-in text log parsers that have common configurations, as described in the "Built-In OPN Configuration Files" section of Parsing Input Text Log Files. However, if there is no configuration file in this list that can adequately parse your text log, you will need to create an OPN configuration file to extend Message Analyzer's parsing capabilities.

More Information

To learn more about OPN configuration files, see the latter sections of Parsing Input Text Log Files.
To learn more about how to create an OPN configuration file, download the OPN Configuration Guide for Text Log Adapter and use it as a development reference to walk through the process of creating a configuration file for your log.

To load saved data from a text log

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.
2. Click the **New Session** button on the Message Analyzer **Start Page** to create a blank session that displays in the **New Session** dialog.
3. In the **New Session** dialog, click the **Files** button under **Add Data Source** to display the configuration features for a Data Retrieval Session.
4. Click the **Add Files** button on the **Files** tab toolbar to display the **Open** dialog, from where you can navigate to one or more text .log files that contain the data you want to view.
5. After you select the .log file/s for which you want to view data, click **Open** to exit the dialog and display the selected .log file/s in the files list on the **Files** tab of the Data Retrieval Session configuration.
6. For each .log file that contains target input data for Message Analyzer, click the **Text Log Configuration** drop-down list to the right of each file and select the appropriate configuration file to parse the log.

   > **NOTE**
   > As described earlier, if there is no configuration file that can parse your log/s, you will need to create a custom OPN configuration file for each log file that has a custom format.

7. Optionally, configure a **Time Filter** if you want to narrow the focus of the data retrieval process to a specific window of time, as described in Applying an Input Time Filter to a Data Retrieval Session.
8. Optionally, select a built-in Filter Expression from the centralized **Library** or configure your own, to create a set of results that focuses on a specific type of data.
9. Ensure that the **Analysis Grid** viewer is selected in the **Start With** drop-down list and click **Start** to exit the **New Session** dialog and begin the data retrieval process.
10. After Message Analyzer has finished loading and parsing the data, observe that the data fields for the text log display appropriately in the **Analysis Grid** viewer.
11. Optionally, add other data fields that were parsed by Message Analyzer to the **Analysis Grid** viewer by using the **Field Chooser Tool Window** to configure them as new columns in the **Analysis Grid**, for enhanced analysis.

Retrieve Data from Log Files in Azure Storage BLOB Containers or from an Azure Table

The procedures in this section show how to retrieve data from one or more logs that are saved in Azure Storage binary large object (BLOB) containers and from an Azure table. Note that when you retrieve data from an Azure log, a **Text Log Configuration** file is required for parsing the log data. However, this is not the case for Azure tables because Message Analyzer automatically parses each property in an Azure table as a field. For more background information about retrieving Azure data, see Handling Azure Data.

To access, load, and view log data from Azure storage BLOB containers

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer**

icon to launch Message Analyzer.

2. Ensure that you have the **AzureStorageLog** parser. You will have it if you have auto-synced the **Azure Storage Parsers Version 1.0** asset collection from the **Asset Manager** dialog and restarted Message Analyzer. If not, follow the general instructions in Using the AzureStorageLog Parser to obtain the parser. The Message Analyzer restart is required after you auto-sync to download the **Azure Storage Parsers** package.

3. Click the global Message Analyzer **File** menu, highlight **Open**, and then select **From Other File Sources** to open the **File Selector** dialog.

4. In the **File Selector** dialog, click the **Add Azure Connection** button to display the **Add Azure Storage Connection** dialog and then specify an **Account name** and **Account key** in the appropriate text boxes — see Accessing Log Data in Azure Storage BLOB Containers.

5. In the Add Azure Storage Connection dialog, ensure that the **Use HTTPS (Recommended)** option is selected as the **Connection** protocol, then click **OK** to exit the dialog.

   A top-level **Azure Storage** node should appear in the left pane of the **File Selector** dialog, along with a data feed subnode beneath it.

6. Expand the feed subnode to expose the **Blobs** container and then navigate to the log/s that contain the data you want to load into Message Analyzer.

7. Select one or more .log files that contain the data of interest and then click **OK** to exit the **File Selector** dialog.

   The **New Session** dialog opens in configuration mode for a Data Retrieval Session, with the Azure .log files that you specified appearing in the files list.

8. In the **Text Log Configuration** drop-down list that displays to the right of each file name, select the **AzureStorageLog** configuration file.

9. Optionally, configure an input **Time Filter** in the **New Session** dialog to narrow the scope of retrieved data to a specified window of time, in order to create a focused data set and accelerate Message Analyzer parsing and loading time.

10. Optionally, if you have downloaded the **Azure Storage Filters** asset collection through the **Asset Manager** dialog, select a built-in Azure Filter Expression from the **Azure Storage** category in the drop-down list of the centralized filter **Library** that appears in the **New Session** dialog, to further focus your retrieval results to a specific type of data.

11. Ensure that the **Analysis Grid** viewer is selected in the **Start With** drop-down list and then click **Start** to exit the **New Session** dialog.

    Message Analyzer begins to parse and load the Azure .log data into the **Analysis Grid** viewer.

12. When data retrieval is complete, observe that the .log data displays in the **Analysis Grid** viewer with data populating the **MessageNumber**, **Timestamp**, **Module**, and **Summary** columns for each message by default.

    Given that the default view **Layout** for Azure .log data is minimalistic, you might want to specify one of several other Azure **Layouts** for the **Analysis Grid** viewer by selecting it from the **Layouts** drop-down list on the **Analysis Grid** viewer toolbar. You can also populate any existing **Layout** with additional columns of data, based on Azure .log field names, as described in the next few steps.

13. Open the **Field Chooser Tool Window**, if it is not already open, by selecting **Field Chooser** from the **Windows** submenu of the Message Analyzer **Tools** menu.

14. In **Field Chooser**, navigate to the **AzureStorageLog** node and then expand it to display the **AzureStorageLogEntry** node. In turn, when you expand this node, **Field Chooser** will display all the fields that Message Analyzer parsed with the use of the **AzureStorageLog** configuration file.

15. Add a new data column to the **Analysis Grid** viewer for a field of interest by selecting a particular field name and clicking the **Add** button in the **Field Chooser** window. You can also add a new column by right-clicking the field and then selecting the **Add As Column** command in the context menu that displays.

    Perform this step for as many field-columns that you want to add for analysis purposes.

    > **TIP**
    > You also have the option to add columns to the **Analysis Grid** by right-clicking a field in the **Details Tool Window** and selecting the **Add Column for < *fieldname* >** command, where *fieldname* is a placeholder for an actual field name in **Details**. Note that you can also use the right-click method for fields in **Details** to add a view **Filter** or to apply grouping to isolate different field values into Groups for enhanced analysis.

16. Alternatively, if you have downloaded the **Azure Storage View Layouts** asset collection through use of the **Asset Manager** dialog, change the **Analysis Grid** viewer column **Layout** by selecting the **Storage Log** item in the **Layout** drop-down list on the **Analysis Grid** toolbar. This **Layout** is specifically designed to include key data fields for analysis of Azure storage logs, such as **ClientRequestId**, **EndToEndLatency**, **ServerLatency**, **RequestStatus**, **StatusCode**, and so on.

    In addition, if you have downloaded the **Azure Storage Charts** asset collection through use of the **Asset Manager** dialog, you can view latency statistics for your Azure storage log data by displaying the **Azure Storage End2End Latency** and **Azure Storage Server Latency Layouts** for the **Chart** viewer, which are accessible from the **Charts** drop-down of the **New Viewer** drop-down list on the global Message Analyzer toolbar.

More Information
To learn more about working with Azure logs, see Retrieving Azure Storage Blob Data.
To learn more about the **Asset Manager**, see Managing Message Analyzer Assets.
To learn more about using the **Field Chooser**, see the Field Chooser Tool Window topic.
To learn more about view **Filters**, see Applying and Managing Filters.
To learn more about the **Analysis Grid** grouping feature, see Using the Analysis Grid Group Feature.
To learn more about creating grouped data views, see the Grouping Viewer topic.

To access, load, and view data stored in an Azure table

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.

2. On the Message Analyzer **Start Page**, click the **New Session** button to create a blank session that displays in the **New Session** dialog.

3. In the **New Session** dialog, click **Azure** under **Add Data Source** to open a Data Retrieval Session from where you can target Azure Storage table data.

4. On the **Azure** tab of the **New Session** dialog, enter the **Account Name**, **Account Key**, and **Table Name** connection information in the appropriate text boxes. See Retrieving Azure Storage Table Data to locate this information.

5. Under the text boxes, select the appropriate connection **Protocol** option to use, which will be either **HTTP** or **HTTPS**.

6. In the **Start With** drop-down list, ensure that the **Analysis Grid** viewer is selected.

7. Begin the data retrieval process by clicking the **Start** button in the **New Session** dialog.
8. When data loading completes, observe that the Azure table data displays in the **Analysis Grid** viewer, with the field data displaying in the **Summary** column.
9. Optionally, display selected Azure table fields in separate columns in the **Analysis Grid** viewer, by utilizing the **Field Chooser Tool Window** to expose other data fields of interest. You will need to look for a module in **Field Chooser** that is associated with the Azure table you specified to find the fields you can add for enhanced data analysis.

See Also

Retrieving Azure Storage Blob Data
Retrieving Azure Storage Table Data

Configuring Session Scenarios with Selected Data Sources

19 minutes to read

When you are creating a Live Trace Session or a Data Retrieval Session with Message Analyzer, you can make use of a flexible session framework that enables you to create session scenarios based on single or multiple data sources. For example, you might want to load data into Message Analyzer from a single set of related files, or you might be interested in correlating data from a particular system Event Log with a related set of trace results.

This section provides guidelines on how to use Message Analyzer to acquire data from one or more supported input data sources in any combination that you wish. You can take advantage of this capability when you need to correlate and analyze related data from different sources, given the enhanced analysis perspective that you can obtain in this context. This section describes input data source configurations, the data sources that are supported by Message Analyzer, an overview of creating input loading and capture configurations, along with a summary of guidelines that you can follow when configuring Live Trace Sessions or Data Retrieval Sessions with single or multiple input data sources.

As indicated in Starting a Message Analyzer Session, all session configuration begins with the **New Session** dialog.

> **IMPORTANT**
> Message Analyzer does not support running a Live Trace Session and a Data Retrieval Session at the same time. Moreover, if you choose to run a Live Trace Session, all other supported input data sources are unavailable to add to the session.

The material of this section is covered in the following topics:

Input Data Source Configurations
Supported Data Sources
Live Trace Scenario Configuration
Data Retrieval Scenario Configuration
Correlating Data from Multiple Input Sources

Input Data Source Configurations

When configuring session scenarios, you have the option to utilize a single input data source or multiple input data sources of different types. Drilling down further, you can also create session scenarios that use multiple inputs from one particular data source type that you choose. For example, you might create a session scenario that uses either of the following data source configurations:

- A single supported data source, optionally with multiple instances of such a single source, where each instance specifies a unique session configuration with different filtering.

  A single source might consist of a live trace that targets one or more computers on which to capture live data, or imported data that is sourced from a single file, a set of specified files, and/or multiple instances of either of these with different filtering configurations applied to each.

- Two or more different supported data sources, optionally with multiple instances of any particular

source, where each instance specifies a unique session configuration with different filtering.

Different data sources might consist of two separate sets of files along with a PowerShell query, a SQL query, data from a specified Event Log, and/or multiple instances of these with different filtering configurations applied, that is, where it makes sense to do so.

The figure that follows is a Data Retrieval Session that is configured to retrieve data from two file sources and a PowerShell query. It also shows the other data source types that Message Analyzer supports on the **New Data Source** tab. Any of these sources could be added to this data loading configuration if correlation of data from such sources made sense.

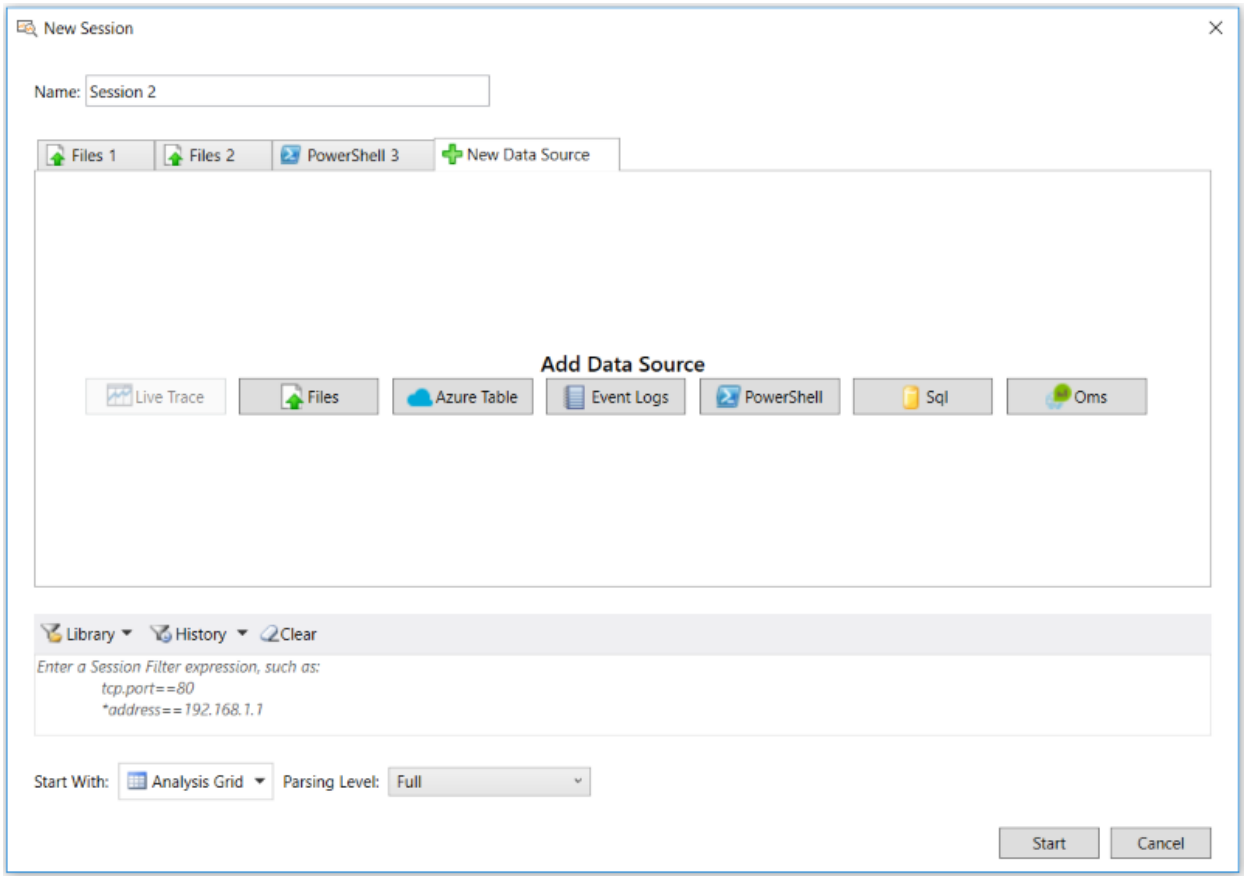

Figure 39: Session data source configuration options

Live Capture Configurations

For Live Trace Sessions, this means that you can create one or more unique capture configurations to acquire live data from the local computer and/or from one or more remote computers. To create this type of session scenario, you would click the **Live Trace** button under **Add Data Source** in the **New Session** dialog. In the **New Session** dialog, you can accept the default and collect live data from the local computer, or you can specify one or more remote computers on which to capture live data. If you want to create another capture configuration to acquire live data, for example from a different set of computers, simply click the **New Data Source** tab and then click the **Live Trace** button again to display the Live Trace Session configuration controls on another **Live Trace** tab. On each tab, the inherent controls enable you to create unique capture configurations that use different built-in **Trace Scenarios**, different ETW providers and provider configurations, and different types of filtering that together can create a capture focus that will return specific data from specified computers for analysis. You can then run all capture configurations simultaneously on the computers that you specified.

More Information

To learn more about creating capture configurations for Message Analyzer, see Configuring a Live Trace Session.

Data Loading Configurations

For Data Retrieval Sessions, this means that you can create a unique data loading configuration that specifies a single data source consisting of a set of saved files that contain the data you want to examine with Message Analyzer. To create this type of session scenario, you would click the **Files** button under **Add Data Source** in the **New Session** dialog. The controls for configuring a Data Retrieval Session then appear on a new **Files** tab, which can include various types of filtering features, truncation, and text log parsers. However, as previously indicated, you have the option to create multiple *sets* of input files, each with different loading configurations, by simply clicking the **New Data Source** tab and then clicking the **Files** button again. You could also create unique loading configurations that specify different target data sources. For example, you could combine **Files**, **Event Log**, **PowerShell** query results, and data from other sources that you can load into Message Analyzer all at the same time for correlation and analysis. In addition, you have the option to specify multiple inputs for any of these sources, for example, multiple PowerShell queries that you specify on different **PowerShell** data source tabs.

More Information

To learn more about creating data loading configurations for Message Analyzer, see Configuring a Data Retrieval Session.

Supported Data Sources

As previously indicated, Message Analyzer provides a separate configuration space for each type of data source, which appears as a separate tab in the **New Session** dialog when you click a particular data source button under **Add Data Source** in the dialog. The input data sources that Message Analyzer supports consist of the following:

> **IMPORTANT**
>
> Some of the following are preview features. To use a preview feature, ensure that it is selected in the **Preview Features** list on the **Features** tab of the **Options** dialog, which you can access from the global Message Analyzer **Tools** menu. If you enable a previously disabled feature, you will need to restart Message Analyzer in order to use the feature.

- **Live Trace** — enables you to capture live data from the network. You can create the input configuration for a Live Trace Session by clicking the **Live Trace** button under **Add Data Source** in the **New Session** dialog. A summary of the configuration features for a Live Trace Session is indicated below in Live Trace Scenario Configuration.
- **Files** — enables you to retrieve data from saved trace and log files. You can create the input configuration for a Data Retrieval Session by clicking the **Files** button under **Add Data Source** in the **New Session** dialog. A summary of the configuration features for a Live Trace Session is indicated below in Data Retrieval Scenario Configuration.
- **Azure Table** — enables you to access data from an Azure table. You can create the input configuration by clicking the **Azure Table** button under **Add Data Source** in the **New Session** dialog.

  You can also acquire input data from Azure storage binary large object (BLOB) logs with the use of the **File Selector**, which is accessible from the global Message Analyzer **File** menu, by selecting **Open** and then clicking the **From Other File Sources** command.

  The configuration features that you can utilize when targeting Azure data is described in Handling Azure Data.
- **Event logs** — a preview feature that enables you to access data from **Microsoft Event Viewer** logs, such as **Applications and Services**, **Windows**, and others. You can create the input configuration by clicking the **Event Logs** button under **Add Data Source** in the **New Session** dialog.

The configuration features that you can utilize when targeting Event Log data is described in Loading System Event Log Data.

- **PowerShell** query — a preview feature that enables you to access data that is output by a PowerShell query, which you create with one or more PowerShell cmdlets. You can create the input configuration by clicking the **PowerShell** button under **Add Data Source** in the **New Session** dialog.

  The configuration features that you can utilize when targeting PowerShell query output data is described in Deriving Input Data with PowerShell Scripts.

- **SQL** query — a preview feature that enables you to access data from a SQL database table. You can create the input configuration by clicking the **Sql** button under **Add Data Source** in the **New Session** dialog.

  The configuration features that you can utilize when targeting SQL table data is described in Loading SQL Data.

- **WPP-Generated Events** — Message Analyzer can process Windows software trace preprocessor (WPP)-generated events. Because WPP events make use of the ETW framework, Message Analyzer can capture them live or load them from a saved event trace log (ETL) file.

- **Operations Management Suite (OMS) logs** — Message Analyzer can load OMS log data, which enables you to leverage Message Analyzer data viewers and analysis capabilities when working with this type of data. To facilitate the process, Message Analyzer provides a search interface to OMS Log Analytics that you can access through the **Oms** data source feature of the **New Session** dialog during Data Retrieval Session configuration.

Live Trace Scenario Configuration

To start the configuration for a Live Trace Session, open the **New Session** dialog from the **Start Page** and click the **Live Trace** button to display the configuration features for capturing data live with Message Analyzer. The configuration controls are contained on a **Live Trace** tab, which enable you to specify **Target Computers**, a built-in **Trace Scenario**, additional ETW providers, filtering configurations, and so on. Note that whenever you click the **New Data Source** tab in the **New Session** dialog for a Live Trace Session, you have the option to open another **Live Trace** tab in which you can specify a different **Trace Scenario** along with specified target hosts, unique filtering criteria, and so on. In this session scenario, you can optionally specify any combination of the following for each selected **Trace Scenario**:

- **Computers** — the target computers on which to capture data.

- **Providers** — additional system ETW providers to enhance the scope of data capture, optionally with **Keyword** and/or **Level** filters set.

- **Provider Filters** — filtering configurations in the **Advanced Settings** dialog for message providers such as the **Microsoft-PEF-NDIS-PacketCapture** provider, **Microsoft-Windows-NDIS-PacketCapture** provider, and the **Microsoft-PEF-WFP-MessageProvider**. The **Advanced Settings** dialog is accessible by clicking the provider's **Configure** link in the **ETW Providers** list on any **Live Trace** tab, after you select a **Trace Scenario** that contains one of the indicated providers. This enables you to create a unique filtering configuration for the selected **Trace Scenario** on each **Live Trace** tab.

- **ETW Session configuration** — settings for the underlying ETW Session for each **Trace Scenario**, which are accessible by clicking the **Configure ETW Session** button on each corresponding **Live Trace** tab.

- **Session Filtering** — a **Session Filter** will apply to any specified **Trace Scenario**.

- **Customized Parsing** — a **Parsing Level** will apply to any specified **Trace Scenario**.

After you start the Live Trace Session, Message Analyzer applies the filtering configurations and other settings that you specified to the live capture on the target computer/s. If you specified any remote computers in the **Target Computers** list on any **Live Trace** tab, Message Analyzer creates a separate sub-session of the Live Trace Session that captures messages on each specified remote computer. The data from all sub-sessions is returned and aggregated to the originating Live Trace Session when the session is stopped.

> **TIP**
>
> After you have captured data from multiple data sources, in many cases you can organize them into separate data source groups in the **Analysis Grid** viewer by right-clicking the **DataSource** field in the **General** category of **Field Chooser** and then selecting the **Add as Grouping** command in the context menu that appears.

Live Trace Scenario Guidelines Summary

The following scenarios provide guidelines that summarize the approaches you can take to session configuration when you are preparing to capture messages live with Message Analyzer:

- **Local Capture : Single Live Trace configuration tab as data source** — enables you to capture messages from the local host with a single **Trace Scenario** and message provider configuration specified on one **Live Trace** tab when running a single Live Trace Session. Filtering configurations expectedly apply to the local host on which data is being captured.
- **Local Capture : Multiple Live Trace configuration tabs as data sources** — enables you to capture messages from the local host with different **Trace Scenarios** and message provider configurations specified on separate **Live Trace** tabs when running a single Live Trace Session. Enables you to focus the trace results on messages at one or more stack levels or messages that pass specific filtering criteria.
- **Remote Capture : Single Live Trace configuration tab as data source** — enables you to capture messages concurrently from different target hosts with a *single* **Trace Scenario** and message provider configuration specified on one **Live Trace** tab when running a single Live Trace Session. Filtering configurations apply equally to all hosts on which data is being captured. Each trace that is running on a specified remote host is considered a sub-session of the originating Live Trace Session. Data that is captured in each sub-session is aggregated into the overall Live Trace Session results that display in a chosen data viewer.
- **Remote Capture : Multiple Live Trace configuration tabs as data sources** — enables you to capture messages concurrently from different target hosts with *different* **Trace Scenarios** and provider configurations specified on two or more **Live Trace** tabs when running a single Live Trace Session. You can apply unique filtering configurations independently for the providers in each **Trace Scenario** on different **Live Trace** tabs.

  For example, you might specify different **Advanced Settings** for the **Microsoft-Windows-NDIS-PacketCapture** provider on each **Live Trace** tab. You could also utilize a different message provider and settings on the second or a subsequent **Live Trace** tab, such as the **Microsoft-PEF-WFP-MessageProvider**. Each trace that is running on a specified remote host is considered a sub-session of the originating Live Trace Session. Data that is captured in each sub-session is aggregated into the overall Live Trace Session results that display in a chosen data viewer.

> **NOTE**
>
> In the **Target Computers** list on each **Live Trace** tab that you select during Live Trace Session configuration, you can specify a different set of hosts on which to capture data.

Data Retrieval Scenario Configuration

To start the configuration for a Data Retrieval Session, open the **New Session** dialog from the **Start Page** and click any button under **Add Data Sources** to display the configuration features for loading the particular type of data into Message Analyzer that you are targeting. For example, if you clicked the **Files** button, then you can use the **Add Files** feature in the **New Session** dialog to locate and target input files that contain the message data you want to load. After you specify an initial data source by clicking an appropriate button under **Add Data Source**, you can click the **New Data Source** tab in the **New Session** dialog to specify another data source type or another of the same source type to add to your initial input data source configuration.

For example, if you initially clicked the **Files** tab, you can specify a set of target input files from which to retrieve data by using the **Add Files** feature. At this point you can add a different data source as described in Supported Data Sources, or you can create another set of input files by clicking the **New Data Source** tab, clicking the **Files** button, and then locating new files with the **Add Files** feature. In this example session scenario, you can optionally apply a differently configured input **Time Filter** to each target set of input files. After you start the Data Retrieval Session, Message Analyzer filters, loads, and chronologically merges all the data into a single viewer, which by default is the **Analysis Grid** viewer.

In any Data Retrieval Session, you can specify settings for the features that follow, although some of these might not have an effect on all the supported input data source types from which you can select. For example, in a Data Retrieval Session that uses the **Files** button to import data from saved trace and log files, you can specify an input **Time Filter** to limit the data you retrieve to a specified window of time. Note that a **Time Filter** is available for use only with a target set of input files on the **Files** tab of a Data Retrieval Session, and not for any other data source:

- **Session Filter** — the effects apply globally across all input data sources in a Data Retrieval Session.
- **Parsing Level** — the effects apply globally across all input data sources in a Data Retrieval Session.
- **Time Filter** — the effects apply only to specific input data source types in a Data Retrieval Session.

> **TIP**
> After you have loaded data from multiple data sources, in many cases you can organize them into separate data source groups in the **Analysis Grid** viewer by right-clicking the **DataSource** field in the **General** category of **Field Chooser** and then selecting the **Add as Grouping** command in the context menu that appears.

Data Retrieval Scenario Guidelines

The following scenarios provide guidelines that summarize the approaches you can take to session configuration when you are preparing to load messages from one or more input sources into Message Analyzer:

- **Single data source and configuration tab** — enables you to retrieve messages from a single input data source that is specified on a single configuration tab in the **New Session** dialog. If you are retrieving data from one or more saved files or logs in this scenario, filtering configurations such as an input **Time Filter** or **Session Filter** apply equally to all specified input files. The Data Retrieval Session begins as soon as you click the **Start** button in the **New Session** dialog. The data that you load into Message Analyzer will appear in the viewer that you choose from the **Start With** drop-down list in the **New Session** dialog.

  > **TIP**
  > If the input data source that you choose in this scenario is **Files**, each file that you target for data retrieval is effectively a data source.

- **Multiple data sources and configuration tabs** — enables you to retrieve messages from multiple

input data sources that are each specified on a separate data source configuration tab in the **New Session** dialog. For example, if you are retrieving data from saved files in this scenario and you are using multiple data source tabs with a different set of input files on each tab, you can configure an input **Time Filter** that is specific only to the set of input files that appear on the tab where you configure the filter. This enables you to independently apply different **Time Filter** parameters to different sets of input files on separate data source tabs. However, if you configure a**Session Filter** or **Parsing Level**, these will apply equally to all specified input files on all data source tabs in the session.

By utilizing multiple data source configurations with different filtering, you can create an output that merges specific messages from multiple input sources that you specified on two or more data source tabs. The Data Retrieval Session begins as soon as you click the **Start** button in the **New Session** dialog. The data that you load into Message Analyzer will appear in the viewer that you choose from the **Start With** drop-down list in the **New Session** dialog.

Correlating Data from Multiple Input Sources

Message Analyzer provides tools that are specifically designed to help you deal with the convergence of related data from multiple input sources. For example, if you had a hypothetical trace and log file that referred to an identical command with different names, you could create a **Union** that enables you to display the correlated field values under the **Union** name in a single data column of the **Analysis Grid** viewer. You can do this by using the **Add as Column** command from the **Field Chooser Tool Window** context menu.

You could also make use of the **Group** feature in the **Analysis Grid** viewer to reorganize the data into grouped message configurations based on the values of chosen columns. For example, you could group all messages in a set of trace results by the values in the **ProcessId** column (a Global Property), so that you can view the message volume associated with each process. Moreover, you might also use the **Grouping** viewer, which performs a similar function, although this viewer provides built-in view **Layouts** that you can select for a more focused analysis of specific groups of data. The **Grouping** viewer also enables you to correlate messages that exist in any particular group, which can include a **Data Source** group, by driving the display of those messages in the **Analysis Grid** viewer for a detailed level of analysis. You can also make use of the **Selection Tool Window** to keep track of correlated messages that you select, for example, the messages that you selected in the **Analysis Grid** viewer, or groups of messages that you selected in the **Grouping** viewer when the **Selection Mode** is active.

Other Message Analyzer tools that support data correlation include **Filters**, **Viewpoints**, and sorting. For example, you can use a view **Filter** to isolate specific data that you want to examine, by filtering for only the types of messages you want to analyze. You could also use a **Viewpoint** to remove messages above a selected **Viewpoint** module that are irrelevant to your analysis goals. Sorting can be particularly useful when you have acquired data from multiple sources that have time stamp differences. In many cases, you will want to correlate your input data sources by time so that you can more easily correlate data that is identical or closely related, but with different time stamps. For example, you could sort the **Timestamp** column in the **Analysis Grid** viewer, and when necessary, add a **Time Shift** to messages from a log that gathered data in a different time zone or where message data is offset because of other factors such as skewed system times.

More Information

To learn more about **Unions**, see Configuring and Managing Message Analyzer Unions.
To learn more about the **Group** feature in the **Analysis Grid** viewer, see Using the Analysis Grid Group Feature.
To learn more about the **Grouping** viewer, see the Grouping Viewer topic.
To learn more about the **Selection Tool Window**, see the Selection Tool Window topic.
To learn more about **Filters**, see Applying and Managing Filters.
To learn more about **Viewpoints**, see Applying and Managing Viewpoints.
To learn more about setting time shifts, see Setting Time Shifts.

See Also

Acquiring Data From Other Input Sources
Configuring a Live Trace Session
Configuring a Remote Capture
Configuring a Data Retrieval Session

Editing Existing Sessions

8 minutes to read

After loading data into Message Analyzer or capturing it live, you have the option to reconfigure your original Data Retrieval Session or Live Trace Session, respectively, so that you can alter the session results to create a different analysis perspective. Session reconfiguration enables you to work with your data until you find the messaging context and data presentation format that helps you resolve the issue on which you are working.

Using the Edit Session Dialog

You can modify any existing session from the **Edit Session** dialog, which is accessible by clicking the global Message Analyzer **Session** menu and then clicking the **Edit Session** command, any time after initial session data is displayed in a viewer such as the **Analysis Grid**. You can also click the **Edit Session** button on the global Message Analyzer toolbar to launch the **Edit Session** dialog.

The session configuration that displays in the **Edit Session** dialog depends on the session viewer tab that has focus (below the global Message Analyzer toolbar). If you have one or more data viewers open for a particular session, selecting any viewer tab in that session and then clicking the **Edit Session** button will launch the **Edit Session** dialog and display the initial configuration that you specified for *that* session.

The figure that follows shows the **Edit Session** dialog as it first appears in **Restricted Edit** mode.

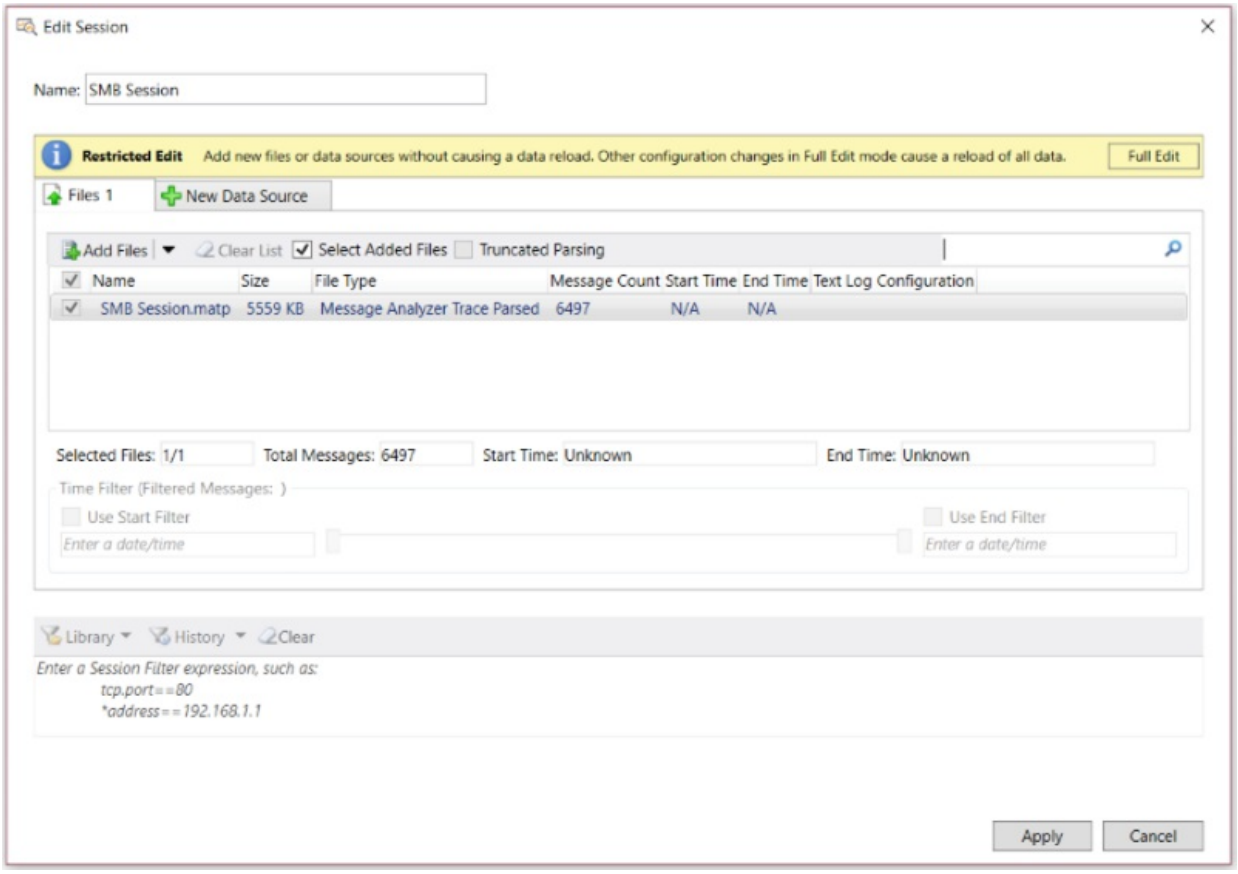

Figure 40: Edit Session dialog in Restricted Edit mode

Editing Modes for Data Retrieval Sessions

If you open the **Edit Session** dialog for a session in which data was loaded into Message Analyzer from a saved trace or log file, the **Edit Session** dialog opens in the **Restricted Edit** mode, which limits the session

changes that you can make to simply adding new files that contain data to be added to the existing displayed message collection. In this mode, Message Analyzer can load the data from new files without having to reload the data from existing files in the files list. To obtain access to all the configuration features for such a session, as described in the Configuring Session Changes, click the **Full Edit** button. However, any configuration changes that you make in this mode will require that Message Analyzer reload all the data, including data from existing files and any new files that you specify. In this case, you might notice that Message Analyzer has a slower loading performance, depending on the size of the new input files.

> **TIP**
> Even if you used the **Open** or **Recent Files** feature to load and display trace or log file data in a Message Analyzer viewer, you still have the option to return to session configuration to make changes. For example, after opening a large trace file through either of these features, you might realize that you need to narrow down the scope of messages to retrieve and then reload the data. You could then define the messages you want to retrieve by applying a **Session Filter** or a **Time Filter** to the data contained in the target input files.

Editing a Live Trace Session

If you recently ran a Live Trace Session but have not yet saved the data, you can return to the initial session configuration by clicking the **Edit Session** button on the global Message Analyzer toolbar, or by clicking the **Edit Session** item in the global Message Analyzer **Session** menu. You might do this if you want to apply various tools and techniques that will refocus the displayed results and create a unique perspective on the existing message collection for analysis purposes, prior to saving the data.

There are no editing restrictions in this scenario, so you have access to all the configuration features described in Configuring Session Changes. After you make changes to the Live Trace Session configuration, simply click the **Apply** button You will then need to restart your session by clicking the **Restart** arrow/button on the global Message Analyzer toolbar to apply the changes you specified. For example, you might have specified other hosts on which to capture message data, chosen a different **Trace Scenario**, added or removed a system ETW provider, added a new filter or modified an existing one, and so on.

> **NOTE**
> If you would rather create and configure an entirely new session, you can do so by clicking either the **Files**, **Live Trace**, **Azure Table**, or other button under **Add Data Source** in the **New Session** dialog, to display the configuration features for a new Data Retrieval Session or Live Trace Session, as appropriate, from where you can name, configure, and start a new session.

Editing a Data Retrieval Session

If you have loaded data from one or more saved files or logs into Message Analyzer, you can make changes from the **Edit Session** dialog and Message Analyzer will apply your changes directly to the existing results display without a reload of existing data, provided that the changes you made were in **Restricted Edit** mode and consisted of adding new files only. When this is the case, the data from the new files will be appended to the existing data in a viewer such as the **Analysis Grid**. Otherwise, any other changes you make such as filtering must be enabled in **Full Edit** mode, in which case, a full reload of all data will occur. At this time, all the changes that you specified will be applied to the reloaded data in this scenario.

Configuring Session Changes

The changes that you can make to an existing session are limited to the type of session you are reconfiguring, as follows:

- **Data Retrieval Session** — you can modify this type of session in the following ways:
  - Add more saved input files to the original files list, to target additional data to analyze.
  - Add more **Files** tabs by using the **New Data Source** tab to target different tabbed sets of input files from which to load data, as described in Configuring Session Scenarios with Selected Data Sources.
  - Specify an input **Time Filter** or reconfigure one that you specified in the initial session configuration.
  - Specify a different **Time Filter** for input files on another **Files** tab, if you added one, as described in Configuring Session Scenarios with Selected Data Sources.
  - Remove the **Session Filter** that you specified in the initial session configuration, reconfigure it, or select a new built-in **Session Filter**. The goal is to apply filtering criteria that is specifically tailored for the data that is being loaded in the updated session.
  - Select a **Parsing Level** that is different from the initial session configuration, or set one for the first time.
  - Select or unselect the **Truncated Parsing** mode, depending on its current setting, if you want to improve performance when loading data from a file that contains truncated (headers only) messages, for example a .cap file.
- **Live Trace Session** — you can modify this type of session in the following ways, that is, if you have not yet saved the session data:
  - Specify different hosts on which to capture data in remote tracing scenarios.
  - Specify a different **Trace Scenario** from the **Select Scenario** drop-down list on the **ETW Providers** toolbar of the **Live Trace** tab.
  - Specify a different **Trace Scenario** on a newly added **Live Trace** tab, as a **New Data Source**, as described in Configuring Session Scenarios with Selected Data Sources.
  - Add one or more system ETW providers to the **ETW Providers** list by specifying them in the **Add System Providers** dialog, which is accessible by clicking the **Add Providers** drop-down list on the **ETW Providers** toolbar.
  - Select (or write) a new **Session Filter**, or modify one that you specified in the original session configuration. Ensure that you are applying filtering criteria that is specifically tailored for the data to be captured in the updated session.
  - Specify or modify additional filters, such as event **Keyword** bitmask and **Level** filters, **Fast Filters**, **WFP Layer Set** filters, **Fast Filter Groups**, NDIS stack or Hyper-V-Switch extension layer packet filters, an adapter filter, and other advanced filters, depending on the message providers in use. For example, you could specify **Advanced Settings** to create provider configurations that enable you to:

    — Capture data directionally on a specified network adapter, for example, in a **Local Network Interfaces Trace Scenario**.

    — Configure a single **Fast Filter** or several logically chained **Fast Filters** for a provider, for example, in **Local Network Interfaces** scenarios on computers running the Windows 7, Windows 8, or Windows Server 2012 operating system.

    — Capture remote traffic on host adapters or on one or more virtual machines (VMs) that are serviced by a Hyper-V-Switch on a remote Windows 8.1, Windows Server 2012 R2, or Windows

10 host, along with specifying NDIS layer and Hyper-V-Switch extension layer packet traversal paths and other special filtering configurations, for example, in **Remote Network Interfaces** scenarios.

- Set a **Parsing Level** or select a different one.
- Modify and optimize the **ETW Session Configuration**, for example, if you think you are dropping packets because of inadequate ETW buffer configuration settings.

More Information

To learn more about selecting a **Trace Scenario**, see Selecting a Trace Scenario.

To learn more about adding and modifying providers, see Adding a System ETW Provider and Modifying Default Provider Settings.

To learn more about how to use a **Session Filter**, see Working with Session Filters in a Live Trace Session and Applying a Session Filter to a Data Retrieval Session.

To learn more about optimizing your ETW session configuration, see Specifying Advanced ETW Session Configuration Settings.

To learn more about **Parsing Levels**, see Setting the Session Parsing Level.

To learn more about how to configure **Advanced Settings** for the **PEF-NDIS** and **Windows-NDIS** providers, see the topics Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog and Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

Viewing Message Data

3 minutes to read

This section describes various aspects of viewing and analyzing data with Message Analyzer's built-in data viewers. Included is conceptual information about the Message Analyzer data viewing infrastructure, the functions of the built-in data viewers, choosing data viewers, manipulating data through viewer and associated features, along with working interactively with other data viewers and **Tool Windows**.

Message Analyzer enables you to evaluate your message data with the primary analysis surface known as the **Analysis Grid** viewer, which provides a rich set of analysis features. Other common data viewers include the **Grouping**, **Pattern Match**, **Gantt**, and other viewers that are described in the Data Viewers section. Message Analyzer also provides a **Chart** viewer for which you can select a host of built-in view **Layouts**. These layouts use graphic visualizer components to provide a high-level overview of message data in various formats or to present focused details and statistical summaries of the data. The built-in **Layouts** for **Charts** are provided by default in every Message Analyzer installation and enable you to uniquely enhance your analysis perspectives for problem solving. Note that you can specify these view **Layouts** for data that you either captured with a Live Trace Session or loaded through a Data Retrieval Session.

Go To Session Analysis Tools
Given that viewing message data and analyzing message data are closely related, you can proceed directly to a brief summary of the analysis tools that are available in Message Analyzer. By reviewing the topic immediately below, you will have a better idea of how to use the Message Analyzer features that are described in this section as analysis tools, which includes data viewers, tool windows, and other features that interact with message data during the analysis process:
Analyzing Message Data

What You Will Learn
In the topics of this section indicated below, you can first review some background information about the Message Analyzer viewing infrastructure and then learn how to use it to present data in various viewer formats, select **Chart** viewer **Layouts** for high-level overviews of data, manipulate the data presented in those viewers, and work with the interactive **Tool Windows**.

In This Section

Data Viewer Concepts — learn about various aspects of the viewing infrastructure, which includes single- and multi-instance data viewers, message- and session-specific **Tool Windows**, data presentation configurations, and the types of data displayed in the Message Analyzer viewing infrastructure.

Data Viewers — review the functions of the built-in data viewers that Message Analyzer provides by default, such as the **Analysis Grid**, **Grouping**, and **Pattern Match** viewers, along with other viewers that are provided as preview features, for example, the **Gantt** and **Interaction** viewers. You can also review the functions of the built-in **Layouts** for the **Chart** viewer.

Session Data Viewer Options — learn how to set a default data viewer for all sessions and how to override this setting in session configuration, and also review the locations from which you can select data viewers.

Common Data Viewer Features — learn how the application of view **Filters**, **Viewpoints**, and **Time Filters** from the Filtering Toolbar enable you to isolate specific data to enhance analysis and problem solving. Also learn about how **Time Shifts** and **Time Format** settings can affect your view of data. In addition, learn how to filter on Data Sources, create and use **Aliases** and **Unions** in the **Analysis Grid** viewer, and also discover how the assets

described in this section integrate with other Message Analyzer analysis features.

Tool Windows — learn how **Tool Windows** interact with data viewers and other windows to display additional message details such as field information, message stack, hexadecimal packet values, diagnosis, and decryption results data. Also learn about the **Tool Windows** that support data analysis capabilities, such as the **Details**, **Diagnostics**, **Message Stack**, **Selection**, and **Field Chooser** windows.

Working with Message Analyzer Window Layouts — learn how to create a custom work environment for the type of troubleshooting and analysis you regularly perform, by choosing a built-in **Window Layout** that organizes data viewers and **Tool Windows** into preset configurations that range from simple to more complex.

Working With Message Analyzer Profiles — learn how to create a focused analysis environment by enabling a built-in or custom-designed **Profile** configured with a data viewer and layout preset that automatically displays whenever you are loading data from a specific type of input file for which the enabled **Profile** is designed, for example a *.cap, *.etl, or *.log file, and so on.

Go To Procedures

To proceed directly to procedures that demonstrate the viewer features described in this section, see **Procedures: Using the Data Viewing Features**.

Data Viewer Concepts

7 minutes to read

Message Analyzer provides a set of default and preview data viewers that enable you to display session results in different viewing formats, whether you are working with a Data Retrieval Session or a Live Trace Session, as described in Starting a Message Analyzer Session. The ability to do so is made possible by the Message Analyzer data viewing infrastructure. The viewing infrastructure enables you to specify the data viewers that provide the unique analysis perspectives that are most advantageous for streamlining your data assessment tasks.

Data Viewing Infrastructure

The Message Analyzer data viewing infrastructure makes use of various data viewers and various types of **Tool Windows** that are either message-specific, session-specific, or provide annotation capabilities. Many data viewers, including the **Analysis Grid** and **Gannt** viewer, are known as *multi-instance* viewers. Multi-instance viewers provide the capability to display message data in multiple instances of the same viewer type. Message-specific windows are typically driven by multi-instance viewers. They reflect data based on the selected messages or fields of an in-focus, multi-instance data viewer. Because message-specific windows can have only a single instance in display, they are considered *single-instance* windows. Multi-instance viewers typically interact with single-instance windows to provide additional data details and presentation enhancements.

For example, by selecting a row of message data in the **Analysis Grid** viewer, the message-specific **Details** and **Message Data** single-instance **Tool Windows** display message field data and hexadecimal field values, respectively, for the selected message. A session-specific window such as **Diagnostics** is driven by session selection. For example, if you have multiple Live Trace Session and Data Retrieval Session viewer tabs displayed, the **Diagnostics** window displays the associated diagnosis summary information for the particular session that corresponds to the viewer tab that you select.

> **NOTE**
> The Message Analyzer viewing infrastructure also enables interaction between some **Tool Windows**. For example, selection of message fields in the single-instance **Details Tool Window** drives interaction with the single-instance **Message Data** and **Field Data Tool Windows**.

Data Presentation Configurations

The Message Analyzer viewing infrastructure represents message data in several presentation configurations, as follows:

- **Tree grid** format — provides a familiar grid view for displaying, filtering, grouping, and analyzing trace or log data, as shown in the Analysis Grid Viewer topic.
- **Graphic visualizers** — consists of several types of graphic displays, as follows:
  - **Chart** viewer layouts — provides graphic data visualizers that are contained in **Layouts** for the **Chart** viewer, which includes built-in and user-configurable **Layouts** that employ **Bar** element, **Pie** slice, **Table** grid, and **Timeline** graphic visualizer components. A built-in **Layout** for the **Chart** viewer that contains a **Bar** element visualizer component is shown in the Chart Viewer Layouts topic.
  - **Gantt** viewer format — this viewer provides an at-a-glance graphic view of message dispersion across a trace timeline that is presented as color-coded protocol module identifiers, with

source/destination address message pairs shown in the y-axis and timestamps in the x-axis. You can see an example of the **Gantt** viewer data visualizer in the Gantt Viewer topic.

  - **Interaction** viewer formats — this viewer contains **Swimlane**, **Mapping**, **Diagram**, and **Chord** viewing formats in which you can view IP conversations between end points in a set of trace results. The **Interaction** viewer is shown with the **Swimlane** configuration in the Interaction Viewer topic.

  - **Message Summary Tiles** format — this viewer uses data tiles to provide a high-level overview of major trace statistics, data summaries, and other important values for any set of trace results. The **Message Summary Tiles** viewer is shown in the Message Summary Tiles Viewer topic.

  - **Message Summary List** format — this viewer uses three data lists to summarize **Module**, **Endpoint**, and **Diagnostic** message count to provide basic information at-a-glance for a set of trace results. The **Message Summary Tiles** viewer is shown in the Message Summary Lists Viewer topic.

- **Other viewing formats** — other unique viewing formats consist of the following:

  - **Grouping** viewer format — organizes traffic into summary hierarchies based on **Grouping** viewer **Layouts** that contain predefined message field Groups in nested configurations. Enables you to expose data at top-level that can normally be difficult to find. The **Grouping** viewer with the **ProcessName and Conversations** view **Layout** is shown in the Grouping Viewer topic.

  - **Pattern Match** viewer format — organizes data into several results panes that display after you execute a built-in Pattern expression that locates a particular message pattern for which the Pattern expression is configured. To see an example of results from executing the **TCP Three-Way Handshake Pattern** expression, see the Pattern Match Viewer topic.

- **Tool data** — provided by message-specific and session-specific **Tool Windows** that serve as analysis tools by working interactively with a data viewer that is in-focus. Tool data provides additional message details and data presentation perspectives that are relative to data selection from an in-focus viewer. Some common **Tool Windows** are shown together with the **Analysis Grid** viewer in the Tool Windows topic.

Data Viewing Infrastructure Implementation

The most common Message Analyzer viewing infrastructure components with which you will typically work the most are described in this section. They display in the main analysis surface where all data viewers appear and includes the **Analysis Grid** viewer, **Tool Windows**, and **Chart** viewer **Layouts**, among others. In Message Analyzer, these viewing components are interactive and integrated such that message selection in one viewing component drives the display of related data such as low-level details or high-level message summaries in one or more other viewing components, providing that integrated components are currently displayed. The following types of message data are implemented in these viewing components:

- **Message detail summaries** — displayed in the tree grid format of the **Analysis Grid** viewer, as described in the Analysis Grid Viewer topic. This format contains session results data that is presented as expandable message nodes in a stacked configuration. Each parent node in the stacked configuration represents an Operation or top-level transaction, while child nodes represent the underlying modules (protocols or providers) that supported such Operations or transactions, including reassembled message fragments that arrived at varying times. In this documentation, these child nodes are also referred to as the *message origins* or the *origins tree*.

- **Message field and value details** — displayed in tabular format and reflect the names, values, types, bit offsets, and bit lengths of the fields of most messages that are selected in the **Analysis Grid** viewer. This information is contained in the message-specific and dockable **Details Tool Window** that displays by default beneath one or more session viewer tabs, as described in the Message Details Tool Window topic.

- **Message field hexadecimal values** — displayed in a message-specific, configurable, and dockable

Message Data Tool Window that shows the hexadecimal values of fields selected in the **Details Tool Window** or messages selected in the **Analysis Grid** viewer. See the Message Data Tool Window topic for further details.

- **Message diagnostics data** — displayed in the session-specific and dockable **Diagnostics Tool Window** that enables quick location of embedded diagnosis message types, summarizes message group counts, and synchronizes diagnosis message selection with their parent messages in various data viewers of a selected session, as described in the Diagnostics Tool Window topic. Note that diagnostic data also displays in the default **DiagnosisTypes** column in the **Analysis Grid** viewer.
- **Message layer data** — displayed in a message-specific and dockable **Message Stack Tool Window** that displays the full stack in several viewing formats for any message selected in the **Analysis Grid** viewer, as described in the Message Stack Tool Window topic. Enables quick exploration of the network architecture for a selected message without manually expanding message nodes in the **Analysis Grid** viewer.
- **Message top-level summary data** — displayed in built-in **Layouts** for the **Chart** viewer, for example, the HTTP Content Type Volumes layout, or in custom, user-designed **Layouts**. The **Layouts** consist of graphical and statistical summary visualizer components that can provide a high-level overview of network traffic, focus on lower-level message details, or display other message activity that occurs across a set of trace results, as described in the Chart Viewer Layouts topic.

  Layouts for the **Chart** viewer condense and encapsulate large volumes of data into visual-graphic presentation formats that are easy to understand at a glance, enabling you to bypass the examination of thousands of messages to obtain a similar level of assessment. This helps to simplify the analysis of data trends, patterns, structures, relationships, and failures, so that you can more efficiently solve messaging problems.

  Note that you can modify the configuration of any built-in **Layout** for the **Chart** viewer and that you can save it under a different name as a new user Library item that you can share with others, as described in Managing User Libraries.

> **NOTE**
> Message Analyzer data viewers are available from the locations described in Session Data Viewer Options.

More Information
To learn more about the multi-instance data viewers, see the Data Viewers topic.
To learn more about the message-specific and session-specific **Tool Windows**, see the Tool Windows topic.
To learn more about creating and managing your own **Layouts** that contain data visualizer components for the **Chart** viewer, see Extending Message Analyzer Data Viewing Capabilities.

Data Viewers

6 minutes to read

In any Analysis Session, Message Analyzer enables you to use data viewers to present message data in various formats to create unique analysis perspectives. You can choose data viewers during session configuration in the **New Session** dialog prior to starting a session, or afterwards in an Analysis Session where you have displayed session results. Some of these are default viewers, while others are preview feature viewers that you can enable, try out, and provide feedback to Microsoft to initiate possible improvements. Several viewers can utilize numerous **Layouts** that change the default view of data and provide advantages such as the following:

- Display different message field data to target a specific analysis environment.
- Change the presentation format to summarize important data at a high-level and provide immediate insights into possible issues.
- Drive important information to top-level for easy viewing or create a window into hidden but significant data that otherwise might be difficult or laborious to achieve through manual methods in a large data set, if even possible at all.
- Perform calculations that expose interesting statistics, extract key information, or enhance the analysis process by utilizing custom global property and annotation values.

The data viewers, the **Layouts** that apply to specific data viewers, and the graphic visualizer **Layouts** for the **Chart** viewer are all described in this section, along with related functional descriptions, how to launch these data viewers, how to display different viewer **Layouts** in those viewers that employ them, and how to use these components to solve diagnostic problems.

The table that follows describes the default data viewers and the preview feature viewers that Message Analyzer provides and also identifies the viewers that employ **Layouts** along with the asset collection Libraries in which they exist.

Table 11. Message Analyzer Data Viewers and Layouts

VIEWER NAME	DESCRIPTION	LAYOUTS	VIEWER TYPE*	ASSET COLLECTION
Analysis Grid	The primary analysis surface that consists of a default tree grid display containing a default column layout with expandable parent and child message nodes that expose top-level transactions/Operations and message origins, respectively, along with message details, in a tree-view structure.	Yes — **Layouts** drop-down list above Filter text box.	**Default**	**Message Analyzer View Layouts**

VIEWER NAME	DESCRIPTION	LAYOUTS	VIEWER TYPE*	ASSET COLLECTION
Grouping	Enables you to organize your traffic into summary hierarchies based on **Grouping** viewer **Layouts** that contain predefined message field groups. These groups exist in nested configurations that create focus on specific data by extracting it from a set of trace results and exposing it at different levels in the group hierarchy. Enables you to drill down into important data while at the same time interactively correlate the group data with the **Analysis Grid** viewer display through group selection.	Yes — **Layouts** drop-down list on the **Grouping** viewer toolbar.	**Default**	**Message Analyzer Grouping View Layouts**
Pattern Match	Enables you to detect message behaviors, patterns, or repeating value sequences within a message collection, based on execution of a predefined or custom-designed OPN **Pattern** expression. Also provides a summary analysis of **Matches**, **Matched Instances** details, and the correlated **Messages**.	None	**Default**	**Message Analyzer Sequence Expressions**

VIEWER NAME	DESCRIPTION	LAYOUTS	VIEWER TYPE*	ASSET COLLECTION
Gannt	Consists of a graphic visualizer presentation that provides a quick view of message dispersion across a trace timeline that is presented as color-coded protocol module identifiers, with source/destination address message pairs in the y-axis orientation and timestamps in the x-axis orientation.	None	**Preview**	N/A
Chart	Defines a generic category of data viewer that uses different built-in **Layouts** that each present data with a single visualizer component such as a grid, bar element, pie-chart, or timeline component with chart formulas applied to create a targeted analysis context. **Layouts** typically consist of top-level message summaries that provide a high-level overview of the data in a set of trace results.	Yes — **Layouts** drop-down enabled in the **Session** menu, only while a **Chart** is displayed. Also accessible from the **New Viewer** drop-down list.	**Default**	**Message Analyzer Chart View Layouts**

VIEWER NAME	DESCRIPTION	LAYOUTS	VIEWER TYPE*	ASSET COLLECTION
Interaction	A swim lane diagram along with several other viewing formats that provide different graphic presentations of the computer nodes and endpoints that exchanged messages within the time boundaries of a trace. Also correlates the IP conversations that took place for the protocols or modules that participated in message exchanges and provides message summary details in hover-over pop-ups.	None	**Preview**	N/A
Message Summary Tiles	Provides a high-level overview of major trace statistics and important values for quick top-level analysis of results.	None	**Preview**	N/A
Message Summary Lists	Provides a high-level overview of major trace statistics and key data points for quick top-level analysis of results.	None	**Preview**	N/A
PerfMon	Enables you to view data from Microsoft Performance Monitor logs in the *.blg file format. You can view this data in Message Analyzer similarly to the way it appears in Performance Monitor while taking advantage of Message Analyzer's data selection, organization, and analysis capabilities.	None	**Preview**	N/A

VIEWER NAME	DESCRIPTION	LAYOUTS	VIEWER TYPE*	ASSET COLLECTION
Charts (Deprecated)	Contains the old Charts that existed prior to Message Analyzer v1.4. For example, you can still display the **Protocol Dashboard** viewer and others in their original configuration of visualizer components.	None	**Preview**	**Message Analyzer Charts**

*A Default viewer is coded into the Message Analyzer UI and cannot be disabled the way a preview feature can be disabled or enabled.

> **IMPORTANT**
> To use any preview feature, you must enable it on the **Features** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. To enable a feature, place a check mark in the check box next to the feature name and then restart Message Analyzer. To disable a feature, remove the check mark so that it will no longer be accessible after the next Message Analyzer restart, that is, until you re-enable it once again.

> **TIP**
> You can specify any data viewer as the default for all sessions by choosing one in the **Default Profile** pane on the **Profiles** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu.

Chart Viewer Layouts Library

The **Chart** viewer **Layouts** Library exists in the **Message Analyzer Chart View Layout** asset collection and is accessible from the **Chart** drop-down list that displays in the **New Viewer** drop-down list; this list is accessible from the locations specified in Session Data Viewer Options. Each **Chart** viewer **Layout** that exists in the specified Library contains a single data visualizer component that typically provides a top-level data summary that you can view for a quick high-level assessment of data. The data visualizers consist of grid, bar element, pie-chart, and timeline components and can display most types of data that you need to present. The **Layouts** exist in the following categories, which are accessible from the **Manage Chart Layout** dialog only while a **Chart** viewer is displayed. You can locate this dialog from the global Message Analyzer **Session** menu by highlighting **Chart**, **Layout**, and then clicking the **Manage** item in the **Manage Layouts** drop-down list.

- **HTTP**
- **General**
- **Network**
- **Netlogon**
- **Common**
- **File Sharing**

The **Chart** viewer **Layouts** that exist in these categories are described in the Chart Viewer Layouts topic. Note that you can also create your own **Layouts** to contain a data visualizer that you choose and data formulas that you configure, and you can create your own categories in which to place them under the default **My Items** top-

level category in the **Layouts** Library.

More Information

To learn more about how to create **Chart** viewer **Layouts**, see Extending Message Analyzer Data Viewing Capabilities.

To learn more about Message Analyzer assets, see Managing Message Analyzer Assets.

See Also

Chart Viewer Layouts

Analysis Grid Viewer

14 minutes to read

The details of raw message data are presented in the Message Analyzer **Analysis Grid** viewer. This is the viewer that displays by default for all Data Retrieval Sessions and Live Trace Sessions, unless you either change the default **Default Viewer** on the **Profiles** tab of the global **Options** dialog that is accessible from the global Message Analyzer **Tools** menu, or if you simply select a different viewer from the **Start With** drop-down list in the **New Session** dialog prior to starting a session. The **Analysis Grid** viewer consists of a default tree grid display that groups related messages together in expandable, stacked configurations of parent and child message nodes that show Operations and other top-level messages above the underlying capture modules (message stack, also called the "origins" in this Operating Guide) and message fragments that supported such transactions across the time boundaries of a trace. In addition, the message **Details**, **Message Data**, **Message Stack**, and **Field Data Tool Windows** display by default when you specify the **Analysis Grid** as a data viewer. Note that message selection in the **Analysis Grid** drives interaction with these **Tool Windows**. The default configuration of the Analysis Grid viewer is shown in the figure that follows.

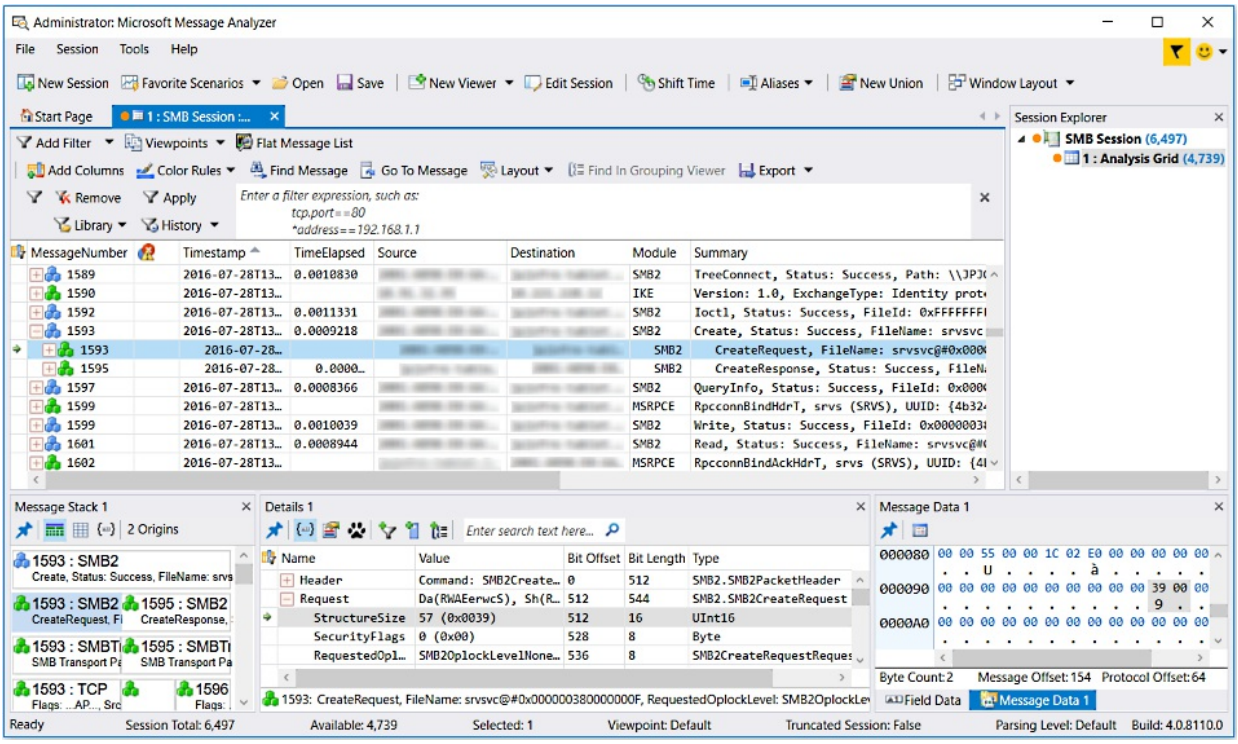

Figure 41: Analysis Grid viewer default configuration

For any message that displays in an **Analysis Grid** message row, you can obtain a quick summary of significant message values at a glance from the data exposed in the default view layout, as described just ahead. You can also use the **Analysis Grid** viewer or the **Tool Windows** that display with it by default, to analyze the network architecture, details, and hexadecimal data for any captured message, which includes the fields and associated values of any protocol or module for which the PEF Runtime parsed data fields.

Note that Message Analyzer now performs resolution of IP addresses to host names, in trace files that contain this information. The results will display in either the **Source** or **Destination** columns of the default **Analysis Grid** view **Layout**.

Default View Layout

By default, the **Analysis Grid** viewer contains the following **Layout** of data columns:

- **MessageNumber** — displays numbers that are assigned to messages that are grouped into top-level messages and operations, which consist of expandable parent message nodes containing child messages that participated in the top-level transaction.
- **Diagnosis** — contains diagnostic icons that you can click to view error details.
- **Timestamp** — displays the time that a message was retrieved.
- **TimeElapsed** — displays a value that is equal to the difference between the **Timestamps** of the last child message and the first child message in the origins tree. This can also represent the amount of time it took for an entire operation to complete.
- **Source** — displays the source address of the highest level message.
- **Destination** — displays the destination address of the highest level message in a transaction.
- **Module** — identifies the protocol or provider message source.
- **Summary** — contains a summary of prominent message information.

> **TIP**
>
> You can use the Message Analyzer **Field Chooser Tool Window** to add a wide array of fields to the default **Analysis Grid** column layout to display additional information about captured or loaded data. If the **Field Chooser** window is not already open, you can display it by right-clicking any column heading in the **Analysis Grid** viewer and then selecting the **Add Columns** command. You can also display it by clicking the **Field Chooser** item in the **Windows** submenu on the Message Analyzer global **Tools** menu.
>
> If you are a Microsoft Network Monitor user and familiarizing yourself with Microsoft Message Analyzer, you might consider adding the **DeltaFromFirst** field to the **Analysis Grid** viewer as a new column to simulate the **Time Offset** field in Network Monitor. Both of these fields in their respective tools provide the sequential running time of each message in a set of trace results. To display the **DeltaFromFirst** field as a column, open the **Field Chooser** and expand the **Global Properties** node. Next, right-click the **DeltaFromFirst** property in the list and select the **Add As Column** command to display it as a new column in the **Analysis Grid** viewer. Note that an instance of the **Analysis Grid** viewer must be in focus when you perform this task.

Tree Grid Features for Data Manipulation

The **Analysis Grid** viewer also provides numerous functions that you can access from the tree grid so that you can manipulate the data display to streamline data assessment. Moreover, these functions enable you to do the following:

- Review a message **Summary** that highlights significant fields and values for any selected message.
- Sort column data to organize and isolate data according to the values of a particular **Analysis Grid** viewer data column.
- Apply a **Group** command to one or more columns to organize data into one or more groups that can provide a quick overview of individual column values or nested groups of values. For example, you might group the default **Source** and **Destination** columns to create an organized display of the message conversations that occurred in a trace.
- Apply a **Column Filter** to any data column to perform a quick search that isolates a particular column value.
- Add other data columns from the **Field Chooser** window to expose the values of specific fields for any module or protocol of interest.
- Display the message **Stack** and message **Field** data inline, for closer scrutiny of the underlying network

layers, fields, and values of selected messages.

- Review error details by clicking any icon in the **DiagnosisTypes** column for a particular message.
- Select items from the right-click context menu of the **Analysis Grid** viewer to perform the following operations:
  - **Show message details** — display message details inline for any selected message by selecting the **Show Details** menu item, to enable detailed analysis of field values.
  - **Create an alias** — create an **Alias** for a column field value, such as a cryptic IPv6 address, by selecting the **Create 'columnEntity' Alias** menu item that displays when you right-click a field value.
  - **Configure and apply a filter** — apply a view **Filter** by choosing the **Add 'columnEntity' to Filter** menu item, to quickly isolate data of interest based on automatically configured filtering criteria associated with the selected column entity.
  - **Filter data in a separate Analysis Grid viewer tab** — select the **Filter 'columnEntity' in New Tab** command to display filtered data in a new instance of the **Analysis Grid** viewer.
  - **Create a Pattern Expression** — select the **Create Pattern** command to open the **Pattern Editor** dialog, from where you can create a pattern expression by configuring a behavior scenario that detects a specified message pattern, as described in Understanding Message Pattern Matching. When you open the **Pattern Editor** from the **Analysis Grid** viewer context menu, the **Quick** tab of the editor is prepopulated with initial information based on the message/s you selected in the **Analysis Grid** viewer prior to opening the editor.
  - **Include hex values** — select the **Include Hex for Numeric Values** command to include a hexadecimal value in parentheses for numeric values that appear in the **Analysis Grid** viewer.
  - **Specify binary value formats** — select the **Display Binary Values As** command to override the corresponding default setting in the **Options** dialog. Selectable values consist of the following, as described in Setting Message Analyzer Global Options:
    - **ASCII**
    - **Hex**
    - **Decimal**
  - **Locate messages in the Grouping viewer** — select the **Find in Grouping Viewer** command to locate a corresponding message that appears in the Grouping Viewer, provided that it is open. This enables you to quickly assess such a message in the context of the Group hierarchy that is created by the currently applied **Grouping** viewer **Layout**. For example, the message you select might appear in a **Network** or **Transport** group in the **Grouping** viewer, enabling you to obtain a deep analysis perspective of the IP conversation or ports that carried the conversation in which the selected message appears.
  - **Display OPN definitions** — select the **Go To 'entityName' Definition** command to display the OPN definition for a module or field that is associated with a selected message.
  - **Add or search for comments** — select the **Add** item in the **Comment** context menu to specify a comment for one or more selected messages, or search for existing comments by selecting the **Find Next** or **Find Previous** items from the **Comment** context menu.
  - **Parse messages on an alternate port** — select the **Parse As** command to open the global **Options** dialog to the **Parsing** tab, from where you can specify a different port for a specified protocol and initiate a reparse of the current set of trace results.

- **Shift Timestamp values** — select the **Shift Time** command to display the **Shift Time** dialog, from where you can specify a shift for the **Timestamp** values of the messages in the current set of trace results.
- **Save specific messages** — select the **Save Selected Messages...** command to save a set of messages that are selected in the **Analysis Grid**. When you select this command, the **Save As** dialog displays to enable you to save the selected messages in a native .matp file in a directory of your choice.
- **Copy data to the clipboard** — select the **Copy Selected Rows** or **Copy 'columnEntity'** command to collect textual message data for record keeping or other purposes.

Analysis Grid Toolbar Features for Data Manipulation

The **Analysis Grid** also has a toolbar that provides several features that you can use to manipulate the data you are displaying, as follows:

- **Add Columns** — add more columns of data from **Field Chooser** that are associated with message types of interest, to expose other field data that does not appear in the **Analysis Grid** viewer default column layout.
- **Color Rules** — specify built-in or user-configurable **Color Rules**, which provide at-a-glance visual message decorations that serve as alerts to invite closer scrutiny of target messages, while minimizing additional diagnostic efforts.
- **Find Messages** — use preset or user-configurable **Find Messages** filtering that enables you to search for specific messages based on filtering criteria.
- **Go To Message** — locate a message based on its message number in one or more data sources.
- **Layout** — specify and manage preset or user-created view **Layouts** that provide data column configurations that are useful for solving common problems or performing repetitive tasks.
- **Find in Grouping Viewer** — locate the Group node within the **Grouping** viewer where a particular message exists that is currently selected in the **Analysis Grid** viewer.
- **Export** — enables you to export either **All** or **Selected** messages in comma separated value (CSV) format. The **Export** command is located on the **Analysis Grid** viewer toolbar.

Analysis Grid Viewer Column Commands

Several commands appear in a context menu that displays when you click any column header in the **Analysis Grid** viewer. These commands enable you to do the following:

- **Group** — enables you to transform the messages displayed in the **Analysis Grid** viewer into a group configuration that filters the data according to criteria that is set by the particular column you right-clicked. For example, if you right-click the **Destination** column and select the **Group** command, the data will be organized into top-level group nodes that are defined by a unique destination IP address where each group contains only those messages that have such a unique address. You might also do the same for the **Source** column to create Source groups nested under the **Destination** groups, in which case you can create a concise organization of the data that shows you the messages related to the IP conversations that took place between specific destination and source computers.

  Note that the **Group** context menu command is not available for the **Summary** column header in the **Analysis Grid** viewer.

More Information

To learn more about the **Analysis Grid** viewer **Group** column command, see [Using the Analysis Grid Group Feature] (using-the-analysis-grid-group-feature.md).

- **Group by Multiple Values** — creates groups based on varying values in a particular field that might be different at various stack levels, for example, IP addresses and Ethernet addresses hidden in the stack. Note that the **Group by Multiple Values** context menu command is not available for the **Summary** column header in the **Analysis Grid** viewer.

More Information
To learn more about the **Analysis Grid** viewer **Group by Multiple Values** column command, see [Using the Analysis Grid Group Feature] (using-the-analysis-grid-group-feature.md).

- **Remove** — enables you to remove any column from the current column **Layout** that is displayed in the **Analysis Grid** viewer. Note that you can restore the default **Layout** for the **Analysis Grid** viewer at any time by selecting the **Restore Application Default Layout** command that displays when you click the **Manage Layouts** menu in the **Layouts** drop-down list on the **Analysis Grid** viewer toolbar.
- **Save as Default User View Layout** — enables you to save any column **Layout** configuration as the user default **Layout**.
- **Load Default User View Layout** — enables you to load the column **Layout** configuration that you saved as the user default **Layout**.
- **Freeze Columns to Left** — enables you to freeze columns to the left of a particular column that you select, such that the horizontal scroll bar starts at the selected column. This feature is particularly useful when you have a column **Layout** that contains many fields, to make it easier to scroll to data. To undo the column freeze, simply apply the **Freeze Columns to Left** command to the **MessageNumber** column.
- **Save Current View Layout As...** — enables you to save the current column **Layout** configuration with a unique **Name**, **Description**, and **Category** specification from the **Edit Item** dialog.
- **Add Columns** — enables you to display the **Field Chooser** window, or to set the focus on it if already displayed.

Other Data Manipulation Feature Locations

Message Analyzer also provides several other locations that contain data manipulation features that can impact the display of data in the **Analysis Grid** viewer, although some of these are also duplicated elsewhere. The other locations for data manipulation features include the Message Analyzer global menus and the global toolbar, which are both located at the top of the Message Analyzer main working interface. Data manipulation commands that are included under global menus consist of the following:

- **Session** global menu — note that several of these commands are duplicated on the global Message Analyzer toolbar:
  - **New Viewer** — enables you to successively select one or more data viewers from this drop-down list to display message data in different presentation formats, for example, the **Grouping**, **Gantt**, **Pattern Match**, and **Chart** viewers.
  - **Edit Session** — enables you to display the **Edit Session** dialog, from where you can modify your Live Trace Session or Data Retrieval Session configuration, which includes different providers, data sources, filters, and so on.

    This command is also duplicated on the Message Analyzer global toolbar.
  - **Reparse** — enables you to reparse the current set of trace results displaying in the **Analysis Grid**

viewer.

- **Shift Time** — displays the **Shift Time** dialog or removes all time shifts that you previously applied.

  This command is also duplicated on the Message Analyzer global toolbar.

- **Data Source Filter** — enables you to apply a filter to the current set of trace results, based on selection of one or more data sources.

- **Active viewer** — displays the items for data manipulation or other functions that are associated with the currently active data viewer. The name of the command that appears in this menu item location changes in response to viewer selection. For example, if the in-focus data viewer is the **Analysis Grid**, the commands on the **Analysis Grid** toolbar appear in a drop-down list in this location. Note that commands from the global Filtering toolbar also appear in this drop-down list no matter which viewer is in focus.

- **Tools** global menu — note that the **Aliases** and **Unions** commands are duplicated on the global Message Analyzer toolbar as drop-down lists.

  - **Windows** — contains all of the selectable **Tool Windows** that Message Analyzer provides for the display of additional data and details that are associated with a set of trace results. These include the **Session Explorer**, **Details**, **Message Data**, **Message Stack**, **Diagnostic**, and **Decryption** windows.

  - **Add-Ins** — contains add-in items such as the **Session Compare Utility**, which performs type checks and field comparisons. Note that the **Compare** utility is currently a preview feature.

  - **Aliases** — enables you to select or unselect a predefined **Alias**, or to manage existing **Aliases** from the **Manage Alias** dialog.

  - **Unions** — enables you to review a list of built-in **Unions** that are included with Message Analyzer. Also enables you to create a new **Union** of your own design by selecting the **New Union** item in the **Unions** submenu to display the **Edit Unions** dialog, from where you can create one. You also have access to the **Manage Union** dialog from this same location.

  - **Asset Manager** — provides management features for common assets that Message Analyzer uses, for example, **Layouts** for the **Chart**, **Grouping**, and **Analysis Grid** viewers. As part of the Message Analyzer sharing infrastructure, **Asset Manager** enables you to download, automatically update, and manage various types of assets, which includes sharing them with others.

  - **Options** — enables you to specify several global options that include default values and selections that can affect Message Analyzer performance, display configurations, and feature activation, as well as the application of **Profiles**, WPP settings, and parsing functionality.

More Information

To learn more about the data manipulation features that you can use when working with the **Analysis Grid** viewer, see the following topics:

Using the Analysis Grid Group Feature
Applying and Managing Analysis Grid Viewer Layouts
Using the Find Message Feature
Using the Go To Message Feature
Filtering Column Data
Using and Managing Color Rules
Applying and Managing Filters
Applying a Time Filter to Session Results

Using the Field Chooser
Setting Time Shifts
Annotation Windows

See Also

Pattern Match Viewer Viewing OPN Source Code](viewing-opn-source-code.md)

Saving Settings

16 minutes to read

When assessing data in an Analysis Session, you will invariably use the many features that Message Analyzer provides for data manipulation, including the default asset collection Library items, such as view **Filters**, **Aliases**, **Viewpoints**, **Charts**, **Color Rules**, viewer **Layouts**, **Pattern Expressions**, and so on, that ship with Message Analyzer. However, it is likely that you will want to create your own customized versions of these Library items or create completely new items of your own to facilitate data analysis. If you create new Library items in the previously described categories for data analysis purposes, Message Analyzer enables you to save your settings for future use and for sharing with others through the Message Analyzer Sharing Infrastructure. When you save an item, it becomes part of the particular asset collection Library with which you are working and such items are available to you from the UI whenever you run Message Analyzer. It is these local Library items that you can share with others by making use of the Message Analyzer Sharing Infrastructure.

There are also additional settings that you can save with session data that include default settings that will apply to all live sessions, and other settings that are the result of manipulating trace data following an Analysis Session, as described in Saving Other Settings.

Saving Settings of Library Items

Message Analyzer provides the facilities for saving the settings of your asset collection Library items from the same location where you apply them for data analysis purposes. You can save the settings of the following asset Library item types in the indicated manner:

- View **Filters** — to save the settings of a custom view **Filter** and add it to your local Filter Expression **Library**, first create one by selecting the **New Filter** item from the **Library** drop-down list on the Filtering toolbar. When you make this selection, the **Edit Filter** dialog displays, from where you can create and save a new Filter Expression based on editing an existing filtering item, or you can create a totally new one by writing a new Filter Expression. You can also specify a filter **Name**, **Description**, and a **Category** to place it in. When you are done with filter configuration, you can save the filter in your local Filter Expression **Library** by clicking the **Save** button in the **Edit Filter** dialog.

> **NOTE**
>
> As indicated in various locations of this operating guide, your local Filter Expression **Library** is a centralized repository that contains the same set of filters that you access when selecting a **Session Filter**, view **Filter**, or **Color Rule** filter. You can also create and save a new Filter Expression in the centralized asset collection **Library** from any of the indicated locations where you select such Filter Expressions.

In addition, if you create and save a Filter Expression that uses an **Alias**, that filter will appear in the centralized Filter Expression **Library** as well. For example, this enables you to use such a Filter Expression containing an **Alias** as a **Session Filter**, **Color Rule** filter, or view **Filter**.

- **Aliases** — to save the settings of a custom **Alias** that you create as a friendly-name substitute for a cryptic field value that might be difficult to work with, and to save it to your **Aliases** collection, first create an **Alias** by right-clicking a field value in the **Analysis Grid** viewer that supports aliasing, such as an IP address, and then select the **Create Alias for <'columnName'>...** item. The **Alias Editor** dialog then displays to enable you to provide the input data required to create a new **Alias**, which consists of **Value**, **Alias** name, **Description**, and **Category** information. When you are done configuring the **Alias**, you can save it in your **Alias** collection by clicking the **Save** button in the **Alias Editor** dialog. After you create an **Alias**, you can

manage it along with other **Alias** collection items from the **Aliases** drop-down list on the Message Analyzer global toolbar or from the **Aliases** submenu on the global Message Analyzer **Tools** menu.

- **Unions** — to save the settings of a custom **Union** that you create to correlate varying field names that have similar values in different data sources, and save it to your **Union** collection, first create a **Union** by clicking the **New Union** button on the global Message Analyzer toolbar or by selecting the **New Union** item from the **Unions** drop-down list in the Message Analyzer **Tools** menu. After you make this selection, the **Edit Union** dialog displays and enables you to provide the input data required to create a **Union**, which consists of **Name**, **Category**, and **Fields** information. You can specify the **Field** information from the **Field Chooser** dialog that displays when you click the **Add** button in the **Edit Union** dialog. When you are done configuring the **Union**, you can save it to your **Union** collection by clicking the **Save** button in the **Edit Union** dialog. After you create a **Union**, you can manage it along with other **Union** collection items from the previously indicated toolbar and locations.

- **Viewpoints** — in the current Message Analyzer release, you can apply **Viewpoints** and manage them, which includes setting **Favorites** along with importing **Viewpoints** that are stored on user-defined shares and creating an export configuration that you can expose to other users through the Message Analyzer Sharing Infrastructure. However, you cannot create any custom **Viewpoints** of your own at this time.

- **Chart Layouts** — to save a new **Layout** for the **Chart** viewer that customizes your data analysis environment, first create one by selecting the **Edit** item in the **Chart** drop-down list that is accessible from the global Message Analyzer **Session** menu. Note that the **Edit** command is available only if you have the **Chart** viewer displayed with a selected **Layout** and in focus. The **Edit** command displays the **Edit Chart Layout** dialog from where you can modify the existing **Chart** configuration as you wish. Thereafter, to save your changes, you must select the **Save Current Layout As** command in the **Layouts** drop-down list that displays after you click the **Layout** item in the **Chart** drop-down list in the **Session** menu. This action adds your new **Layout** to the **Message Analyzer Chart View Layouts** asset collection, which displays in the user Library that appears in the **New Viewer** drop-down list on the global Message Analyzer toolbar, among other places. Note that your new **Layout** is automatically added to the top-level **My Items** category of this asset collection, which you can view and manage from the **Manage Chart Layout** dialog that is accessible from the **Session** menu while a **Chart** is being displayed.

- **Color Rules** — to save the settings of a custom **Color Rule**, the filtering criteria it contains, and to add it to your local **Color Rule** user **Library**, first create one by selecting the **New Color Rule** item from the **Color Rules** drop-down list on the **Analysis Grid** viewer toolbar. When you make this selection, the **Edit Color Rule** dialog displays, from where you can create and save a new **Color Rule**.

- **Layouts** — to save the settings of a custom view **Layout** for the **Analysis Grid** viewer and add it to the local **Layout** user Library, first use the **Field Chooser Tool Window** to add new data columns that will expose the values of message fields that are of critical importance to solving a specific problem. For an example of a built-in column arrangement that focuses on TCP analysis, see the **TCP Deep Packet Analysis with ABSOLUTE Sequence Number Flat** view **Layout**. When you arrive at a useful configuration, you can save it as a new **Layout** that becomes part of the **Message Analyzer View Layouts** asset collection that you can access from your local **Layout** Library on the **Analysis Grid** toolbar, thus enabling you to apply the new **Layout** to the **Analysis Grid** viewer at any time. You can save your new **Layout** by clicking the **Save Current Layout As...** item from the **Layout** drop-down list on the **Analysis Grid** viewer toolbar. This action displays the **Edit Item** dialog, from where you can specify a **Name** and **Description** for the new **Layout** and add it to the default **My Items** category, or you can specify a new **Category**.

  You can also specify any **Layout** you choose to be the default column layout configuration for the **Analysis Grid** viewer, by selecting the **Save Current As Default User Layout** item in the **Manage Layouts** submenu. If you modify your default user **Layout** configuration and you need to restore it to the original version, you can do so by selecting the **Load Default User Layout** item in the same submenu. However, if you need to restore the default **Layout** that ships with Message Analyzer, you can do so by selecting the

Restore Application Default Layout item, which is also a **Manage Layouts** submenu item.

> **NOTE**
> You can also save the settings of **Layouts** that you create for the **Grouping** viewer.

- **Pattern Expressions** — to save the settings of a custom Pattern Expression, first build one by using the features of the **Pattern Editor**, the user interface for which displays when you click the **Create Pattern** button on the **Pattern Match** viewer after you open this viewer from the **New Viewer** drop-down list. When you are done using the **Quick** or **Free Form** tab to create your Pattern Expression, click the **Save Pattern** button on the **Pattern Editor** dialog toolbar to add your new Pattern Expression to the **My Items** category of the **AVAILABLE PATTERNS** list that appears in the **Pattern Match** viewer. This action also adds your new Pattern Expression to the **Message Analyzer Sequence Expressions** asset collection, which comprises a user Library that appears as the indicated patterns list. You also have the option to manage and share Pattern Expressions by selecting the **Manage Patterns** command from the **Pattern Match** drop-down list in the global Message Analyzer **Session** menu.

- **Trace Scenarios** — to save the settings of a custom **Trace Scenario**, first configure one from the **New Session** dialog, which displays when you click the **New Session** button on the Message Analyzer **Start Page**. From the dialog, you can specify the message providers you want to use in addition to the target host/s, filtering, viewer, and parsing level you want to use. You can even customize one of the built-in **Trace Scenarios** to your own design before saving your settings by clicking the **Save Scenario** button on the **ETW Providers** toolbar in the **New Session** dialog. When you save a customized **Trace Scenario**, it is automatically added to the **Message Analyzer Trace Scenarios** asset collection, which appears as the user Library in the **Select Scenario** drop-down list in the **New Session** dialog. Your new **Trace Scenario** and all its settings are stored in the **My Items** category of this drop-down under a subcategory name that you specify. As with all other Message Analyzer asset collections, you can manage and share **Trace Scenario** assets from the **Manage Trace Scenario** dialog, which displays when you click the **Manage Trace Scenarios** command in the **Select Scenario** drop-down list.

Saving Other Settings

Other items that you can save for Analysis Sessions consist of the following. This includes default settings that you specify in the global **Options** dialog that is accessible from the global Message Analyzer **Tools** menu:

- **Comments** — you can add one or more comments to any message that you choose in the **Analysis Grid** viewer. To add comments, you will need to display the **Comments Tool Window**, which is accessible from the **Windows** submenu of the global Message Analyzer **Tools** menu. Thereafter, you can specify a **Title** for each comment, an **Author** name and date-time, and of course the comment text itself. After you configure a comment, you have the option to edit or delete it. When you want to view comments that you previously configured in a set of trace results, you can perform a forward or backward search by clicking the **Next** and **Previous** buttons on the **Comments** tool bar, respectively. As you do this, the messages for which you configured comments are highlighted in the **Analysis Grid** viewer, so you can quickly return to commented messages of interest. Any comments that you specify are saved with the **Analysis Grid** viewer trace results, provided that you save the trace as a native .matp file.

- **Bookmarks** — you can add bookmark settings to any message that you choose in the **Analysis Grid** viewer. You can also add bookmarks to messages in the **Pattern Match** viewer. To add bookmarks, you will need to display the **Bookmarks Tool Window**, which is accessible from the **Windows** submenu of the Message Analyzer **Tools** menu. Thereafter, you can specify a bookmark **Name**, add a **Category** for a bookmark, and configure a color-coded **Flag** that indicates a user-defined condition. By selecting a bookmark message row in the **Bookmarks** window, you can highlight the corresponding messages in the **Analysis Grid** viewer to quickly return to one or more book-marked messages of interest. Any bookmark

settings that you specify are saved with **Analysis Grid** viewer trace results, providing that you save the trace as a native .matp file and that you save all messages. If you save selected messages only or a filtered message set, bookmarks are not saved in the trace file.

- **Time Shifts** — you have the option to save any **Time Shift** settings that you applied to a set of trace results. However, you must save your trace results data in the .matp file format. A **Time Shift** enables you to adjust for machine skew or time zone changes so that you can sync up two or more data sources.
- **Time Display** — you have the option to set the **Time Zone** and **Date and time format** that Message Analyzer will use for all displayed messages in any session where **Timestamps** are used. Formats consist of date and time, or time only. You can specify the **Time Display** option that you want on the **Display** tab of the global **Options** dialog, which is accessible from the Message Analyzer **Tools** menu.
- **Live Trace Message Buffer** — includes options for saving the message buffer **Size** limit and the **Contiguous Drop Percentage** rate for the PEF Runtime, which together specify the rate at which packets are dropped when the live buffer size setting is exceeded. You can change these settings in the **Live Trace Message Buffer** pane on the **General** tab of the **Options** dialog, which is accessible from the Message Analyzer **Tools** menu. If you change the default values for these options, you will automatically save them when you click **OK** to exit the **Options** dialog. These options are global in that they apply to all Live Trace Sessions.
- **Session Viewer Defaults** — Message Analyzer enables you to set the default viewer in which your message data will display, whether you are capturing data in a Live Trace Session or if you are loading data into Message Analyzer through a Data Retrieval Session. The **Default Viewer** drop-down list on the **Profiles** tab of the **Options** dialog enables you to make a default selection from among all data viewers that are available to Message Analyzer. The **Default Viewer** drop-down list contains the same data viewers that exist in the **New Viewer** drop-down list on the global Message Analyzer toolbar and in the **Session Explorer Tool Window** context menu, with the exception of a default **Chart** item (without additional **Layouts**) that exists in the indicated **Default Viewer** drop-down list and the **Start With** drop-down list in the **New Session** dialog. The data viewer that you choose as the default is saved when you click **OK** to exit the **Options** dialog. The selected viewer will then be used by default for any session that you run, unless you specifically change it from the **Start With** drop-down menu of the **New Session** dialog, in which case, the selected viewer applies to the current session only.
- **Text Log Files** — if you select an item from the **Default text log configuration** drop-down list on the **Text Log Files** pane of the **General** tab in the **Options** dialog and save it by clicking **OK** to exit the dialog, it becomes the default configuration file for all text logs from which you load data into Message Analyzer. This setting is saved across Message Analyzer restarts.
- **Binary Values** — by selecting one of the following options on the **Display** tab of the **Options** dialog, you determine the default format in which Message Analyzer displays binary values, for example, in the **Payload** field of the **Details Tool Window**:
  - **Display as ASCII**
  - **Display as Hex**
  - **Display as Decimal**

> **NOTE**
> This setting is persisted across Message Analyzer restarts, just as all settings that you specify in the **Options** dialog are persisted.

- **Decryption Options** — you can add one or more server certificates and passwords in the **Certificates** pane on the **Decryption** tab of the **Options** dialog, to enable Message Analyzer to decrypt traffic that is

encrypted with the Transport Layer Security (TLS) or Secure Sockets Layer (SSL) security protocols. You automatically save these items in the Message Analyzer certificate store when you click the **OK** button to exit the **Options** dialog. Thereafter, you can load a saved trace file or start a Live Trace Session to retrieve the target encrypted traffic and Message Analyzer attempts to decrypt as many messages as possible. Server certificates are saved across Message Analyzer restarts; however, for security purposes, you must manually re-enter passwords after a Message Analyzer restart, prior to decrypting a trace.

- **Session results** — when you save the results of an **Analysis Session**, you have the option to save all session messages, subsets of a message collection that result from the application of various data manipulation items such as **Viewpoints** and view **Filters**. The **Save/Export Session** dialog that is accessible by clicking the **Save** item in the global Message Analyzer **File** menu provides options to save data in this manner, which includes the **All Messages**, **Filtered Messages**, or **Selected Messages** options. For example, a subset of a message collection might reflect the application of a **Time Filter** or view **Filter**, which you would save under the **Filtered Messages** option. You can also save specifically selected messages with the **Save Selected Messages** context menu command in the **Analysis Grid** viewer. You can save your session results in the .matp (parsed) or .cap file format only.

> **NOTE**
> If you save your trace results as a .cap file, **Comments**, **Bookmarks**, and **Time Shift** settings are not saved.

- **Parsing Levels** — if you set a **Parsing Level** when configuring a **New Session** that you run, thereafter when you save the message set that displays in a data viewer such as the **Analysis Grid**, the applied **Parsing Level** and its effects will be reflected in the saved file. Note that you must use the **Open File** feature on the **Start Page** or the **Open** and **From File Explorer** commands on the **File** menu, rather than loading one or more files through a **New Session**, to retrieve the data with the original **Parsing Level** applied.

See Also

Applying and Managing Filters
Using and Managing Message Analyzer Aliases
Configuring and Managing Message Analyzer Unions
Applying and Managing Viewpoints
Using and Managing Color Rules
Managing Chart Viewer Layouts
Applying and Managing Analysis Grid Viewer Layouts
Using the Pattern Editor
Creating and Managing Custom Trace Scenarios
Setting the Session Parsing Level
Parsing Input Text Log Files
Using the Field Chooser
Saving Message Data
Annotation Windows
Setting Time Shifts
Applying a Time Filter to Session Results

Using the Field Chooser

8 minutes to read

The Message Analyzer **Analysis Grid** viewer has a default view **Layout** that contains several columns in which basic message data is displayed; however, the default layout displays only a limited cross-section of the available data. Additional information is also available for many other message fields that you can access by using the **Field Chooser Tool Window**. By adding specific columns, you can expose hidden but important field information that you can examine for greater troubleshooting capabilities with Message Analyzer.

> **NOTE**
>
> The column configuration for the default **Analysis Grid** view **Layout** is described in the Analysis Grid Viewer topic. Also, you can find all the individual columns that are contained in the default layout of the **Analysis Grid** viewer in the **Field Chooser** by using its search facility.

Accessing the Field Chooser

Although the column configuration of the default **Analysis Grid** view **Layout** provides some basic information for any message that you view, there are many more data columns that you can add to the **Analysis Grid** viewer that enable you to focus on specific data fields that contain values for the message types, properties, structures, methods, flags, events, metadata, and so on, of your captured messages. Each of the nodes in the message hierarchies of the **Field Chooser** window use common icons to represent fields, methods, properties, and so on. You can access the **Field Chooser** from the **Analysis Grid** viewer in any of the following ways:

- Click the **Add Columns** icon on the **Analysis Grid** viewer toolbar. This action opens and docks the **Field Chooser** window in its default location. If the **Field Chooser** is already displayed when you click the **Add Columns** icon, then **Field Chooser** simply becomes the active window.
- Click the **Add Columns** command from the **Analysis Grid** submenu in the Message Analyzer global **Session** menu, when the **Analysis Grid** has focus.
- Right-click any **Analysis Grid** column label and then select the **Add Columns...** command in the context menu that appears, which likewise opens and docks the **Field Chooser**, if it is not already open. If **Field Chooser** is already open, the previously indicated action occurs.

> **NOTE**
>
> You can also access the **Field Chooser** from other locations, by doing any of the following:

- Click the **Add Groupings** icon on the **Grouping** viewer toolbar.
- Select the **Field Chooser** item from the **Windows** submenu, which is accessible from the Message Analyzer global **Tools** menu.
- Click the **Field** ellipsis (**...**) in the **Series Fields** pane of the **Edit Chart Layout** dialog or click the **Value** ellipsis in the **Values** pane of the **Edit Chart Layout** dialog, which is accessible by clicking the **Edit** command from the **Chart** drop-down list in the global Message Analyzer **Session** menu whenever a **Chart** has focus.
- Click the **Add** button in the **Edit Union** dialog when you are creating a new **Union** (click the **New Union** button on the Message Analyzer global toolbar; this command is also accessible from the

Message Analyzer global **Tools** menu).

- Click the **Insert Message** button on the **Quick** tab of the **Pattern Editor** dialog, which displays when you click the **Create Pattern** button in the **Pattern Match** viewer. You can also display **Field Chooser** by clicking any ellipsis in the **Criteria** section of the **Quick** tab (click **Insert Criteria**), after you launch the **Pattern Editor** by selecting the **Create Pattern** command from the context menu that displays when you right-click any message in the **Analysis Grid** viewer.

Adding Data Columns to the Analysis Grid with Field Chooser

The **Field Chooser** window contains a top-level tree view of all the message modules and protocols for which Message Analyzer provides parsing based on OPN descriptions. If you expand the nodes of any particular module, you will see the message hierarchy containing the message types, properties, structures, methods, flags, events, or other data fields that are defined for that module or protocol. If you want to view the data for other fields that are relevant to the trace data you captured, you can add a data column for specified fields to the **Analysis Grid** viewer column layout by locating the appropriate module or protocol and the required field names in the **Field Chooser** window, and then doing any of the following while the **Analysis Grid** viewer is in focus:

- Double-click the field name.
- Select the field or other entity that you want to add as a column and then click the **Add** icon in the upper-left corner of the **Field Chooser** window.
- Right-click a field name and select the **Add as Column** item in the context menu.

> **TIP**
> You can also select the **Go to Definition** item in this context menu to open the OPN viewer and highlight the definition of the field you chose.

For example, if you wanted to view the ID of processes for which Message Analyzer captured events, you could add **ProcessId** as a column — found under the **Etw** node or under the **Global Properties** node in the **Field Chooser** — to the **Analysis Grid** viewer. Note that you also have the capability to display a **ProcessName** column from **Global Properties** as well. You are advised that viewing process ID and process names works best when you are displaying data from a *.etl file.

Adding Groups to the Grouping Viewer with Field Chooser

You can also use the **Field Chooser** window to add one or more new Groups to one of the default Grouping **Layouts**, or you can do this when you are creating your own Grouping **Layout**. To add a new Group, you will need to locate one or more fields for a particular message type that you want to add to your **Layout** as a Group. When you locate a field, right-click it and then select the **Add as Grouping** command from the context menu that appears. Note that to add the selected field as a Group, the **Grouping** viewer must be in focus. Otherwise, if the **Analysis Grid** viewer is in focus, selecting the **Add as Grouping** command reorganizes the **Analysis Grid** message display into groups that each contain a set of identical values, where such values are aggregated from all messages that contain those particular values for the field that you added with the **Add as Grouping** command.

More Information
To learn more about adding new Groups to the **Grouping** viewer, see the Grouping Viewer topic.
To learn more about creating groups to enhance your analysis perspectives with the **Analysis Grid** viewer, see Using the Analysis Grid Group Feature.

Searching for Fields

If you want to search for a particular field by name, you can type the name, or a portion thereof, in the search text box at the top of the **Field Chooser** window and search results will display with yellow highlighting, similar to the way Windows Explorer displays search results. After you locate the correct field, you can then add it as a data column to the **Analysis Grid** viewer or as a Group in the **Grouping** viewer, in the previously specified manner. If there is data associated with any field you chose, it will automatically display after you add the field as a new column in the **Analysis Grid** viewer, or as a Group in the **Grouping** viewer.

> **NOTE**
>
> You can remove any **Analysis Grid** column by right-clicking the column label and selecting the **Remove** command in the context menu that displays. You can remove a Group in the **Grouping** viewer by clicking the **x** in the Group label below the **Grouping** viewer toolbar.

> **TIP**
>
> Because the **Field Chooser** window provides a tree-level view of the message hierarchy for each module type that is parsed by the PEF Runtime, it can help you understand how to traverse the message hierarchy to some extent. However, Message Analyzer also provides the Filter IntelliSense service that helps you to discover how to traverse the message hierarchy in a more interactive way; this feature streamlines the task of writing Filter Expressions. Filter IntelliSense is an interactive and intelligent statement completion service that responds to text that you enter in any Filter Expression text box by displaying various elements of the message hierarchies, such as message types, structures, properties, flags, and other fields. The message hierarchies that you traverse with Filter IntelliSense are the same that display in **Field Chooser** window.
>
> Note that you can invoke Filter IntelliSense when configuring Filter Expressions in the **Session Filter** text box in the **New Session** dialog, or in the **Filter** and **Viewpoint Filter** text boxes on the Filtering toolbar that is located just below any session viewer tab.

Adding Analysis Grid Columns from the Details Tool Window

In addition to using the **Field Chooser** window to add more data columns to the current column configuration of the **Analysis Grid** viewer, you can also quickly add a new data column to the **Analysis Grid** viewer based on field names in the **Details Tool Window**, which displays beneath the **Analysis Grid** viewer by default. To do this, right-click any row of field data in the **Details** window and select the **Add 'fieldName' as Column** menu item, where *fieldName* is a placeholder for the actual field, property, or other name in the **Name** column of the **Details** window. Message Analyzer then automatically populates the new named column with values based on parsed data from the currently displayed message collection.

More Information

To learn more about the **Field Chooser**, see the Field Chooser Tool Window topic.
To learn more about Filter IntelliSense, see the Filter IntelliSense Service.
To learn more about Filter Expressions, see Writing Filter Expressions.
To learn more about the **Grouping** viewer, see the Grouping Viewer topic.

See Also

Applying and Managing Analysis Grid Viewer Layouts

Using and Managing Color Rules

7 minutes to read

Message Analyzer provides an important analysis feature known as **Color Rules**, to enable you to define special filters that use color, text, and font styles to decorate and highlight messages that contain specific information in a trace. With this feature, you can apply visual indicators on top of individual message rows in a set of trace results or log data displaying in the **Analysis Grid** viewer, to easily identify specific types of traffic at a glance, thereby lessening the need for additional actions and diagnostics. Rather than as a mechanism that isolates particular messages that meet specific criteria and removes those that do not, such as you have with a view **Filter**, you can think of **Color Rules** as a feature that serves the following purposes:

- **Alerts** — you can use **Color Rules** to serve as a warning or reminder function that invites closer scrutiny and further investigation of certain messages that meet rules criteria.
- **Diagnostics** — you can configure multiple **Color Rules** with varying decoration features to provide an instant visual cue of messages that meet the criteria of multiple rules, based on specified filtering configurations that target diagnostic issues.
- **Organization** — you can organize the view of network/protocol infrastructure at a high level.

With **Color Rules** applied, you can quickly expose subtleties that are not readily obvious in the default display style of the **Analysis Grid** viewer — for example, network direction, message payloads, and error conditions. This can help you effectively evaluate common issues with accuracy and speed. Moreover, you can design multiple **Color Rules** with differing decoration configurations and filtering criteria that apply to specifically related message types, so that the default display style will alter to include the decoration schemes of all the matching **Color Rules**. Using **Color Rules** in this manner can provide a more visually robust indication of some state, condition, or value that you are trying to pinpoint at-a-glance.

For example, you could very quickly expose messages with a particular transport and network protocol in a trace by applying the default **TCP** left gradient and **IPv4 Right Gradient** rules from the **Color Rules** drop-down list on the **Analysis Grid** toolbar. Thereafter, all **TCP** messages that use **IPv4** will display in the **Analysis Grid** viewer with the opposite facing gradient color configurations that are characterized by these rules. Note that the **Analysis Grid** viewer will expose these color configurations by showing them in top-level messages, even if one of those **Color Rules** applies to an unexposed origins message only.

Using the Default Color Rule Asset Collection Library

To help you get started and familiarized with real-world usage examples, Message Analyzer provides a default set of **Color Rules** in several different **Categories** in the **Message Analyzer Color Rules** asset collection **Library**. These categories provide some basic filtering rules, decoration schemes, and text styles that are useful for exposing messages from common protocols such as TCP, ARP, SMB, HTTP, RPC, and so on. Message Analyzer provides the default set of **Color Rules** in the following categories:

- **Network**
- **Azure Storage**
- **Event Log**
- **File Sharing and Authentication**
- **Internet**

- **Event Tracing for Windows**
- **Hardware and Parsing Errors**
- **RPC**

You can select any or all of the default **Color Rules** in your rule Library to immediately apply the associated rule functionality to a currently displaying message set. As a result, the default display color and style of messages in the **Analysis Grid** viewer that match the filtering criteria of the applied **Color Rules** will be overlaid with the decoration scheme and styles of such **Color Rules**. This temporary alteration to the default message display style provides an instant visual cue that enables you to rapidly discover message data that may be of interest in your analysis process. When you deselect one or more **Color Rules**, the associated display schemes and styles are immediately removed from the corresponding messages in the **Analysis Grid** viewer.

Using Gradient Decoration Configurations

You might note that some of the default **Color Rules** in the **Network** category make use of *color gradients*, for example left-to-right and right-to-left gradient configurations. Message Analyzer enables you to specify a gradient background for your **Color Rule** configurations as one way to design visual decoration cues that are distinguishable from each other when messages meet the criteria of multiple rules that overlap. For example, you might design two **Color Rules**, each with different but related filtering rule criteria, or you could apply one of the default **Color Rules** that are configured to produce multiple visual cues when a message meets the criteria of multiple rules.

Several default **Color Rules** in the **Network** category will overlap each other when their rules are met, to provide visual indications of the network and transport layer protocols contained in certain messages in a set of trace results. The **Color Rules** that enable this to occur consist of the **IPv4**, **IPv6**, **TCP**, and **UDP** rules. These rules are configured to work in pairs, as follows:

- **IPv4** and **TCP**
- **IPv4** and **UDP**
- **IPv6** and **TCP**
- **IPv6** and **UDP**

The **Color Rule** filtering criteria in this example is simple. The **IPv4 Color Rule** uses an "IPv4" filter expression; the **TCP Color Rule** uses a "TCP" filter expression; the **UDP Color Rule** uses a "UDP" filter expression; and so on for **IPv6**. When a particular message meets the filtering criteria of any of the **Color Rule** pairs indicated in the example, both directional gradients will be overlaid on the message and rendered easily visible, thus providing a quick visual cue of the particular message types that the specified filtering criteria reflects. For example, in the case of the default **IPv4** and **TCP** rules, **IPv4** uses a **Right-Hand Side** gradient decoration and **TCP** uses a **Left-Hand Side** gradient decoration, such that both decoration schemes will be easily recognized in a message row when the filtering criteria of both these rules is met.

> **NOTE**
> The default gradient **Color Rule** pairs are specifically designed to visually coordinate the display of their gradients in opposite directions when overlaid on a single message that matches the filtering criteria of both rules. This is possible because each **Color Rule** makes appropriate use of a white background for either the **Left-Hand Side** or **Right-Hand Side** color in the **Gradient Background** configuration, so that the overlaid gradients can show through.

Expanding the Default Color Rule Asset Collection Library

If the default **Color Rules** do not provide the features that you need, you can obtain additional **Color Rules** in the

following ways:

- **Download** — download **Message Analyzer Color Rule** asset collection updates from the Message Analyzer feed in the **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu.
- **Import** — use the **Manage Color Rule** dialog to **Import** additional **Color Rule** items directly from a file share or other location where team members have posted such items for sharing.
- **Develop** — create, configure, and save your own **Color Rules** in your local **Color Rule** Library.

Note that you can also **Export** your **Color Rule** items through the Message Analyzer Sharing Infrastructure to a file share that you can configure as a feed to which others may subscribe. In a future Message Analyzer release, the Sharing Infrastructure publishing features will automatically enable others to synchronize to any updates that you make to your **Color Rule** items on the feed. However, you can currently perform a manual configuration process that enables users to synchronize to your item updates, as described in Manual Item Update Synchronization.

> **TIP**
> Because it can be time consuming to create your own **Color Rules**, you should take advantage of **Message Analyzer Color Rule** asset collection updates that are periodically available from Microsoft through the **Message Analyzer** feed in the **Asset Manager** dialog, or from other users that have created and shared their **Color Rules**. To do this, you will make use of the auto-sync and download features for user Library asset collections and the Message Analyzer Sharing Infrastructure, which are available from the **Asset Manager**. Since these features provide a convenient and consistent method for managing and sharing these items, you can quickly learn from others how to apply different methods and approaches to problem diagnosis, and thereby realize improvements in your effectiveness and efficiency as a message analysis specialist.

More Information
To learn more about **Color Rule** management features, see Managing Color Rules.
To learn more about the Message Analyzer Sharing Infrastructure, see the Sharing Infrastructure topic.
To learn more about auto-syncing and downloading user Library asset collection updates, see Managing Asset Collection Downloads and Updates.

See Also

Applying Color Rules

Creating and Modifying Color Rules

5 minutes to read

Message Analyzer enables you to create new **Color Rules** of your own design or you can modify existing ones. Message Analyzer provides an **Example Color Rules** item under **My Items** in the **Color Rules** drop-down list on the **Analysis Grid** viewer toolbar that you can modify however you want for practice purposes, so you can begin familiarizing yourself with **Color Rule** features and functions. After you create a new **Color Rule** or modify an existing one, you can add it to an existing **Category** or populate it to a new **Category** that you define for your local **Message Analyzer Color Rules** asset collection Library. You can perform these operations by using the **Edit Color Rule** dialog, which is accessible in each of the following ways:

- By selecting the **New Color Rule** item from the **Color Rules** drop-down list on the **Analysis Grid** viewer toolbar. Note that this drop-down list is also available from the global Message Analyzer **Session** menu, but only when an **Analysis Grid** viewer tab has focus.
- By right-clicking a default **Color Rule** in a **Category** of your local rule Library and selecting the **Create a Copy** item in the context menu that displays.
- By right-clicking a **Color Rule** in a **Category** from within the **Manage Color Rules** dialog. This dialog is accessible by selecting the **Manage Color Rules** item from the **Color Rules** drop-down list in the previously specified locations.

> **NOTE**
> If you configure a new **Color Rule** from the **Manage Color Rules** dialog by creating a copy of an existing rule and modifying it, or if you directly modify an existing rule, the changes will be added to your local **Color Rule** asset collection Library.

Caution

You cannot use the **Edit Color Rule** dialog to modify any of the built-in **Color Rules** that are provided by default with every Message Analyzer installation.

Creating New Color Rules

To create a new **Color Rule** with the **Edit Color Rule** and to specify a **Category** in which to place it, you will need to do the following:

- Open the **Edit Color Rule** dialog by selecting the **New Color Rule** item from the **Color Rules** drop-down list on the **Analysis Grid** viewer toolbar.
- Enter a name for the new **Color Rule** in the **Name** text box.
- Optionally describe the purpose of the **Color Rule** in the **Description** text box.
- Specify the **Category** that you want to contain the new rule, either by selecting an existing item from the drop-down menu of the **Category** combo box, or by typing a new category name in the **Category** combo box.
- Specify a Filter Expression for the new **Color Rule**.

  To do this, you can manually configure a filter by entering the Filter Expression text directly in the filter text box below the **Edit Color Rule** dialog toolbar, or you can select a built-in Filter Expression from the centralized **Message Analyzer Filters** asset collection **Library**. If you want to manually configure a Filter

Expression, the Filter IntelliSense service is available to provide statement completion assistance.

> **NOTE**
> Whether you specify a manually configured Filter Expression or you modify one of the predefined expressions from the centralized **Library**, Message Analyzer automatically validates that the expression properly compiles when you attempt to **Save** the **Color Rule**. If the filter expression does not compile, you will be unable to save the new **Color Rule**.

- Configure the decoration scheme for the new **Color Rule** by using the controls in the **Style**, **Weight**, **Lines**, and **Colors** panes of the **Edit Color Rule** dialog. As you modify the **Color Rule** decoration scheme, the configuration is displayed in the **Preview** pane at the bottom of the **Edit Color Rule** dialog.

  For each **Color Rule** you create, be sure to specify a unique decoration and text scheme that will readily convey its status to you as a warning, reminder, or diagnostic alert. **Color Rule** decoration schemes consist of a combination of the following visual components:

  - **Foreground** (text), **Gradient Background**, and **Background** colors.
  - **Normal**, **Italic**, and **Oblique** text styles.
  - **Normal** and **Bold** font styles.
  - **Normal**, **Underline**, **Strikethrough**, and **Overline** line text decorations.

- Create the new **Color Rule** and **Category** (if you specified one) by clicking the **Save** button in the **Edit Color Rule** dialog.

> **NOTE**
> After you save a new **Color Rule** or an edited one, it is automatically applied to the current message collection displaying in the **Analysis Grid** viewer.

Modifying Color Rules

Message Analyzer enables you to modify, delete, or create a copy of any **Color Rule** under the **My Items** node only, in the **Color Rules** drop-down list. This includes the **Example Color Rule** that is provided for practice purposes. You can also **Create a Copy** of any predefined **Color Rule**, modify it, and then **Save** it, but you cannot directly modify a predefined **Color Rule** that is provided by default with Message Analyzer. The commands to perform these operations are available as the following context menu items, which display as appropriate for the category node in which you right-click a **Color Rule**:

- **Edit** — available for **Color Rules** under the **My Items** node only. Displays the **Edit Color Rule** dialog, from where you have access to the **Color Rule** configuration described in Creating New Color Rules.
- **Create a Copy** — available for **Color Rules** under the **Message Analyzer** and **My Items** category nodes. Displays the **Edit Color Rule** dialog, from where you can copy an existing **Color Rule**, for example, into a different **Category**. You can also reconfigure the copy as a new rule and place it in a **Category** of choice, which is tantamount to creating a new rule.
- **Delete** — available for **Color Rules** under the **My Items** node only. Removes a single selected **Color Rule** under the **My Items** node.

When you modify a **Color Rule**, for example, to specify a different decoration scheme or filtering configuration, you are using the same **Edit Color Rule** dialog with which you create new **Color Rules**.

TIP

It is advisable to create a backup of your local **Color Rule** Library asset collection, so that you can always restore any custom rules you may have deleted under the **My Items** category of your **Message Analyzer Color Rules** asset collection Library.

More Information

To learn more about managing **Color Rule** items, including sharing them with others, see Managing Color Rules.

To learn more about Filter Expressions, see Writing Filter Expressions and the Filter IntelliSense Service topic.

Applying Color Rules

2 minutes to read

In the Message Analyzer **Analysis Grid** viewer, each row of the message grid displays in a default style. When a message meets the filtering criteria of a **Color Rule**, the colors and styles of the rule that differ from the default row style are applied on top of the associated message row. If the filtering criteria of another **Color Rule** also applies, both styles can be applied to the particular message row in the **Analysis Grid** viewer. For an example of this, see Using Gradient Decoration Configurations.

Applying Color Rules to a Trace

When you are ready to apply **Color Rules** to the results of a Live Trace Session or a Data Retrieval Session, you can select the **Color Rules** that you want to apply in a few different ways, as follows:

- **Individually** — apply specific **Color Rules** by selecting them one at a time.
- **Categorically** — apply all the **Color Rules** in a **Category** by selecting the **Category**, for example, the **Network** category or **My Items** category.
- **Globally** — apply all the **Color Rules** in the local **Message Analyzer Color Rules** asset collection Library by selecting the **Message Analyzer** and **My Items** nodes in the **Color Rules** drop-down list, which is accessible on the **Analysis Grid** viewer toolbar.

Each rule **Category** contains a tree view of **Color Rules** that you can selectively enable or disable by placing a check mark in the check box of a rule, or by removing it, respectively. A **Category** that has only a subset of rules enabled is shown with its check box in the shaded tri-state condition. A **Category** that has all its rules enabled is simply shown with a check mark in its selection check box.

> **NOTE**
>
> If you leave a **Color Rule** or one or more **Categories** of rules enabled and you close the current Live Trace Session or Data Retrieval Session, the selected **Color Rules** will continue to be applied to all subsequent session results that you display in the **Analysis Grid** viewer, until such time that you disable them. Enabled **Color Rules** are also applied across all current Message Analyzer sessions.

See Also

Creating and Modifying Color Rules

Managing Color Rules

6 minutes to read

Message Analyzer enables you to share your local **Message Analyzer Color Rules** asset collection Library items with others, either by employing a user-configured feed or by sharing asset collection items directly with other users on a designated file share. You can share your entire Library as a set, or you can create a subset that includes specific **Color Rules** while excluding others. You can also modify the filtering and decoration scheme for **Color Rules** that you will export and share with others; however, keep in mind that any edits you make to custom **Color Rules** that you intend to export, will modify the corresponding local asset collection Library items. You can view the **Color Rule** items that are available in your local asset collection Library by clicking the **Color Rules** drop-down list on the **Analysis Grid** viewer toolbar.

To create a **Color Rule** export configuration, you must select the **Manage Color Rules** item from the **Color Rules** drop-down list on the previously indicated toolbar to display the **Manage Color Rules** dialog. Note that you can use the **Manage Color Rules** dialog to export and import **Color Rules** directly to and from a user-designated location, respectively, without employing a user-configured feed in the Message Analyzer Sharing Infrastructure.

Managing Color Rule Items

The **Manage Color Rules** dialog provides the following UI elements to enable you to share **Color Rules** with others:

- **Import** — enables you to navigate to a location designated by a user or team that previously posted a **Color Rule** asset collection, so that you can retrieve it and add the **Color Rules** it contains to your local **Library**.

Caution

You should not use the **Import** feature to retrieve **Color Rule** items from a file share that is configured as a feed in the Message Analyzer Sharing Infrastructure. If you use the **Import** feature to retrieve **Color Rule** items from such a user file share, all rules that you select in the **Select Items to Import** dialog will be *added to* your local Library, while any items that have an identical name will not be overwritten, resulting in duplicate rules in your local Library. For **Color Rule** downloads or updates through the Message Analyzer Sharing Infrastructure, you should always click the status icon on the **Downloads** tab of the **Asset Manager** dialog to obtain the download or update options for the **Message Analyzer Color Rules** or **Azure Storage Color Rules** asset collections. For more information about downloading asset collections, see Managing Asset Collection Downloads and Updates. For a procedure that provides an example of importing Library items, see Share Local Library Items on a File Share.

- **Export** — you can use this feature to post a set of **Color Rules** to a designated file share that you intend to configure as a feed in the Message Analyzer Sharing Infrastructure. However, you can also use this feature to post a **Color Rule** asset collection to a file share that is not configured as a user feed in the Message Analyzer Sharing Infrastructure. Note that if you use the **Export** feature to publish *updates* to a **Color Rule** asset collection on a user feed, some manual configuration is necessary for this to be successful, as described in Manual Item Update Synchronization.

- **Delete** — enables you to remove any **Color Rule** in the **My Items Category** of your local **Color Rules** asset collection Library.

  Note that you cannot delete any of the predefined **Color Rules** in your local Library, although you can **Create a Copy** to save in the **My Items Category**. However, to ensure that you do not lose any of your custom-created **Color Rules**, you should save a backup to a designated location so that you can always reimport your rule set as necessary.

- **Rule selection** — enables you to specify the **Color Rules** that you want to include in your export configuration by selecting check box nodes, as follows:
  - **My Items** — includes all **Color Rules** under the **My Items** node in the export configuration.
  - **Message Analyzer** — includes all **Color Rules** in each **Category** node under **Message Analyzer** in the export configuration.
  - **Category** — includes all the **Color Rules** of a specified **Category**, for example, under the **Network** node, in the export configuration.
  - **Color Rules** — includes one or more selected **Color Rules** in your export configuration.
- **Context menu** — the context menu that displays when you right-click a **Color Rule** in the **My Items** category contains the following items:
  - **Edit** — displays the **Edit Color Rule** dialog, from where you can modify any **Color Rule** under the **My Items** node only. If you modify the **Color Rule** configuration or **Category** placement from this location, the changes will apply to your local asset collection Library after you click **Save** in the **Edit Color Rule** dialog.
  - **Create a Copy** — displays the **Edit Color Rule** dialog, from where you can copy an existing **Color Rule**, for example, into a different **Category** in your local asset collection Library. After you make a copy of an existing **Color Rule**, you can reconfigure it as required and move it to a new or existing category under **My Items**; thereafter, your local Library will reflect the changes that you specified.
  - **Delete** — removes a single item from your local asset collection **Library** that you select under the **My Items** node.

    > **IMPORTANT**
    > The right-click **Delete** command removes a single selected item from your local asset collection Library under the **My Items** node only. You can also perform a **Delete** operation from the **Manage Color Rules** dialog by clicking the **Delete** button on the dialog toolbar after selecting one or more **Color Rules**. If you attempt to **Delete** any of the predefined **Color Rules** from this location, you will be prompted for a deletion of all **Color Rules** in the **My Items** category.

> **NOTE**
> When you open the **Edit Color Rules** dialog from the **Manage Color Rules** dialog, by right-clicking a **Color Rule** and selecting either the **Edit** or **Create a Copy** command from the context menu, you will be using the same dialog that displays when you select the **New Color Rule** item in the **Color Rules** drop-down list on the **Analysis Grid** viewer toolbar. The only difference is that when you select the **Edit** or **Create a Copy** command as indicated, you can only *revise* the **Color Rule** configurations, rather than create any new rules. Please note that any modifications you make to a **Color Rule** from the **Edit Color Rule** dialog—when accessing such dialog from the **Manage Color Rules** dialog—including deleting it, changing the decoration scheme, or modifying the **Category**, **Name**, filtering criteria, and so on, will be reflected in your local Library.

Updating the Color Rules Asset Collection

Microsoft provides a default **Message Analyzer** feed on the **Downloads** tab of the **Asset Manager** dialog that enables you to download **Message Analyzer Color Rules** or **Azure Storage Color Rules** asset collections from a Microsoft web service and to synchronize with asset collection updates that are periodically pushed out by the service. At any time, you can perform a download of an auto-synced collection from the **Settings** tab of the **Asset Manager** dialog.

More Information

To learn more about the common **Manage <*AssetType*>** dialog, see Managing User Libraries.

See Also

Creating and Modifying Color Rules

Downloading Assets and Auto-Syncing Updates

Using the Find Message Feature

2 minutes to read

When you are analyzing message data that is displayed in an **Analysis Grid** viewer tab, you have access to the **Find** function to search for and locate individual messages that meet the filtering criteria of a Filter Expression that you select or define. To enable the **Find** filtering mode, click the **Find Message** icon on the **Analysis Grid** viewer toolbar. After you select a predefined Filter Expression from the **Library**, or write one in the text box of the **Find Messages** window, click the **Find** binoculars icon to locate the next message that meets the specified filtering criteria. You can also click the **Find Previous** icon to return to a previous message in the trace results that meets the specified filtering criteria.

Utilizing the Advantages of the Find Message Function

One of the advantages of using the **Find** function to locate messages that meet specific filtering criteria is that you can locate single messages in a large data set relatively quickly. An even more important advantage is the context that the feature provides. Often, the messages that precede or trail a target message can provide an indication of what caused the condition that is reflected by the target message — for example, an error — which can be key to the analysis process. This can also include factors such as message types, field values, time stamps, and so on. By contrast, when you apply a view **Filter**, all messages are filtered out except those that meet the filtering criteria, which can obscure the surrounding context.

Specifying a Filter Expression for the Find Message Function

To specify a Filter Expression for the **Find** function, you can manually configure one or select a built-in expression from the centralized Filter Expression **Library** that is accessible in the **Find Messages** window, just as you would do when specifying a **Session Filter**, view **Filter**, or **Color Rule Filter**. The Filter Expressions that are available for the **Find** function are the identical **Library** items that you can access for **Session Filter** or view **Filter** configuration, and include any custom Filter Expressions that you create.

> **NOTE**
>
> You can use the **Find** function to locate messages in the **Grouping** viewer that meet specified filtering criteria, providing that the **Grouping** viewer is in focus when you click the **Find** binoculars icon.

More Information

To learn more about view **Filters**, see Applying and Managing Filters.

To learn more about **Session Filters**, see Working with Session Filters in a Live Trace Session.

To learn more about sharing Filter Expression Library items with others, see the Sharing Infrastructure topic.

Using the Go To Message Feature

4 minutes to read

To simplify the search for a specific message by message number, Message Analyzer provides the **Go To Message** dialog, which is accessible from the **Analysis Grid** viewer toolbar. As a tool that enables you to quickly navigate to a particular message in a large data set, it can accelerate your analysis process, for example, in analysis scenarios where you are being directed to quickly locate specific data. Note that the **Go To Message** feature is enabled for multiple data sources, meaning that you can elect to search all data sources or a specified source.

Locating Messages

The **Go To Message** feature locates messages in the **Analysis Grid** viewer only, based on entering a numerical message number in the **Go To Message** dialog. This also includes locating messages in any Grouped message display that you create in the **Analysis Grid** viewer. If a match is found, the message of interest is highlighted in the **Analysis Grid** viewer. Even if the message is hidden within unexposed origin tree nodes, the **Go To Message** feature expands the origins nodes to whatever level is necessary to select the message of interest and make that selection visible. In addition, if you happen to have a Grouped display of message data in the **Analysis Grid** viewer, the **Go To Message** feature will expand any Group node in which the message exists, select the message, and thereby make the selection visible. This provides a quick and convenient method for locating specific messages that you want to examine.

Searching Multiple Data Sources

The **Go To Message** dialog also provides several options for searching through trace results, based on selection of sources that collected data for the particular session in which you are searching for messages. These options consist of the following:

- **Search All Data Sources** — select this check box in the **Go To Message** dialog to search all session data sources for a specified message number entry. Currently, the **Go To Message** will find only the first message in any data source that matches the message number you entered in the dialog.
- **Data Source** drop-down list — select an item from this list to specify the data source that contains the message source you want to search for a specified message number entry.

The previously described options appear in the **Go To Message** dialog only if a particular session contains more than one data source. This scenario is common when you have multiple traces or logs that contain related message data that you need to consolidate for analysis purposes. When you load this type of input data configuration into Message Analyzer, it results in messages from such consolidated data sources displaying in chronological order, interlacing them as the chronology requires it. Note that there are two ways that you can create a session that has more than one data source, as follows:

- **Add files** — when configuring a Data Retrieval Session, you can use the **Add Files** feature in the **New Session** dialog to add multiple files to the file list on the **Files** tab. Each file that you add is considered a data source, which could be a saved trace or log file. Thereafter, when you launch the **Go To Message** dialog, the data sources listed in the **Data Source** drop-down list might be identified by a trace file name, log file name, or could be associated with a session GUID.
- **New data sources** — when configuring a Live Trace Session, you can select different **Trace Scenarios** on two or more **Live Trace** tabs to create multiple data sources. To do this, you will create at least one additional **Live Trace** tab to contain a **Trace Scenario** by clicking the **New Data Source** tab.

NOTE

When configuring a Data Retrieval Session, you can likewise make use of the **New Data Source** tab to specify multiple **Files** tabs in the **New Session** dialog, where each **Files** tab contains a different data source, such as a saved trace or log file. On each **Files** tab, you can locate the saved traces or logs that will be your data sources by clicking the **Add Files** button and navigating to those sources.

Displaying the Go To Message Dialog

To display the **Go To Message** dialog, an **Analysis Grid** viewer session tab must be in focus in order to expose the toolbar on which the **Go To Message** button exists. Then, you can either click the **Go To Message** button on the **Analysis Grid** viewer toolbar, or you can simply use the keyboard shortcut `Ctrl+G` to open the dialog.

Improved Method for Locating Messages

Prior to the introduction of this feature, you may have experienced the inconvenience of having to write a **Find Message** filter to locate a message. If you did, then it probably looked similar to the following: `#messagenumber==1234`. Matches found in these cases are located to the top-level message only, which was then highlighted.

At that point, if you performed a subsequent click of the **Find** binoculars icon, a **Find Message** information bar displayed with the message "**Finished finding all messages**", as if the message was found, when actually it was not. When this occurred, you needed to perform multiple successive clicks to expand nodes and navigate the message layers until the actual message of interest was found. Even if the top-level message node and underlying child nodes were completely expanded to expose the message stack, you were still required to perform subsequent clicks of the **Find** binoculars icon to step through the open layers in succession. Even then, the search may have unexpected results at times.

The **Go To Message** feature enables you to avoid all this, as it provides a more convenient and quicker method for locating messages. However, please note that the **Find Messages** feature is still very useful in finding messages that meet the criteria of a **Find Messages** filter that you specify from the **Message Analyzer Filters** asset collection **Library** drop-down list and apply to a search.

Applying and Managing Analysis Grid Viewer Layouts

19 minutes to read

Message Analyzer enables you to apply built-in and custom column layout configurations to the **Analysis Grid** viewer to create unique data viewing capabilities that expose other data fields of interest beyond the default **Analysis Grid** viewer column configuration. These layout configurations are included in the **Message Analyzer View Layouts** asset collection Library, which is accessible from the **Layout** drop-down list on the **Analysis Grid** viewer toolbar. By default, this Library contains a list of built-in view **Layouts** and a customizable **Example** layout, which are included in every Message Analyzer installation. The **Layout** drop-down list also has several commands that enable you to save, restore, and manage the **Layouts**. The main reason why you might change the default **Layout** is to reconfigure the **Analysis Grid** viewer data column arrangement to one that is more applicable to the type of troubleshooting or analysis tasks you are performing. When you select a built-in **Layout**, Message Analyzer adds new data columns that expose additional message field values. By doing this, you can expose hidden data that might be critically important to your data analysis perspectives.

You can change the default column arrangement by applying any built-in **Layout** or by creating, saving, and applying any custom **Layout** of your own. You can create your own custom **Layouts** by using the **Field Chooser Tool Window** to add new columns to the **Analysis Grid** viewer, or you can similarly customize the **Example Layout** however you want. Note that you can even apply the **Group** function to one or more **Analysis Grid** viewer data columns, by right-clicking chosen data columns and selecting the **Group** command from the context menu, and then save the resulting grouped configuration as a new view **Layout** asset collection Library item. You can also sort any column and the sort configuration will be persisted when you save the **Layout**.

The commands that facilitate these features, along with the built-in **Layout** functions, are described in the sections that immediately follow.

Applying Built-in View Layouts

Message Analyzer provides several built-in view **Layouts** that you can quickly apply to a message collection in the **Analysis Grid** viewer by simply selecting them from a drop-down list. These layouts consist of many different column configurations that expose data that can be useful for troubleshooting common problems, for example, when the TCP or HTTP protocol is the focus of analysis. In addition, the built-in **Layouts** provide various grouped message presentations that streamline analysis of data from Perfmon, WPP, ProcMon, and other logs. The built-in **Layouts** are contained in the following subcategories of the top-level **Message Analyzer** category:

- **Azure Storage** category
  - **All .Net Client Columns** — displays all data columns from a .Net client log.
  - **Grouped by ClientRequestId and Module** — provides a grouped **Layout** of **ClientRequestId** at top-level and **Module** as a nested group, for Client, Storage, and Network logs.
  - **Storage Log** — provides a **Layout** that displays key data columns from Azure Storage logs such as **ClientRequestId**, **EndToEndLatencyMS**, **ServerLatencyMS**, **RequestStartTime**, **RequestStatus**, and **StatusCode**.

- **Cluster** category
  - **Cluster Log** — provides a **Layout** that is useful for analyzing Cluster logs with data columns such as **ProcessId**, **InfoLevel**, **Subcomponent**, and **RemainingText**. Enables you to correlate process IDs with various debug levels and Cluster Service components where errors and warnings may have occurred, respectively. This **Layout** is intended to work with the **Cluster Logs Layout** of the **Grouping** viewer to create an interactive analysis environment. In addition, this **Analysis Grid** viewer **Layout** displays by default if you are loading data from a Cluster.log file and the **Cluster Logs Profile** is enabled in the **Options** dialog, as described in the Working With Message Analyzer Profiles topic.

    See the **Cluster Logs Profile** in the table of the indicated topic for more information about how to analyze these logs, where such analysis is enhanced by complementary **Layouts** for the **Grouping** and **Chart** viewers, which are also described there.
- **Windows Event Tracing** category
  - **ETW** — provides a **Layout** that is useful for analyzing ETW messages that contain **ProcessId** and **ThreadId** data. The **ProcessId** is a number that is used by the operating system kernel to uniquely identify an active process for which an ETW provider or some other component is generating events. The **ThreadId** is a unique identifier of an execution thread that is running under a particular process. This **Layout** also includes a **TimeDelta** column that exposes the running time at which each message was captured, similar to the way **TimeOffset** does in the Network Monitor view.

    This **Layout** is intended to work with the **ETW Guids and IDs Layout** of the **Grouping** viewer to create an interactive analysis environment. In addition, this **Analysis Grid** viewer **Layout** automatically displays if data from an event trace log (*.etl) file is loaded into Message Analyzer while the **ETW Analysis Profile** is enabled in the **Options** dialog, as described in the Working With Message Analyzer Profiles topic.

    See the **ETW Analysis Profile** in the table of the indicated topic for more information about how to analyze these logs, where such analysis is enhanced by complementary **Layouts** for the **Grouping** and **Chart** viewers, which is also described there.
- **Examples** category
  - **Event Log (.evtx)** — provides a **Layout** that is useful for analyzing ETW event data. This **Layout** is intended to work with the **Event Viewer Layout** of the **Grouping** viewer to create an interactive analysis environment. In addition, this **Analysis Grid** viewer **Layout** displays by default if you are loading data from a *.evtx file and the **Event Logs Profile** is enabled in the **Options** dialog, as described in the Working With Message Analyzer Profiles topic.

    See the **Event Logs Profile** in the table of the indicated topic for more information about how to analyze these logs, where such analysis is enhanced by complementary **Layouts** for the **Grouping** and **Chart** viewers, which is also described there.
- **HTTP** category
  - **Fiddler SAZ** — provides a **Layout** that is useful for analyzing HTTP data in an environment that is similar to Fiddler, with data fields such as **StatusCode**, **Uri.Schema**, **Method**, **Uri.Host**, **Uri.AbsPath**, **PayloadLength**, **ContentType**, and **Payload**. This **Layout** is intended to work with the **Fiddler Grouping Layout** of the **Grouping** viewer to create an interactive analysis environment. In addition, this **Analysis Grid** viewer **Layout** displays by default if you are loading data from a *.saz file and the **Fiddler Traces Profile** is enabled in the **Options** dialog, as described in the Working With Message Analyzer Profiles topic.

See the **Fiddler Traces Profile** in the table of the indicated topic for more information about how to analyze the data in these files, where such analysis is enhanced by complementary **Layouts** for the **Grouping** and **Chart** viewers, which is also described there.

- **File Sharing** category
  - **File Sharing Perf SMB2/SMB** — provides a **Layout** that enables you to analyze SMB/SMB2 performance in terms of **ResponseTime** and **TimeElapsed** values for SMB/SMB2 request and response messages. Also provides other data columns that are useful for SMB analysis such as **SessionIdName**, **TreeIdNameReference**, **FileNameReference**, and **MessageId**. This **Layout** is intended to work with the **File Sharing SMB/SMB2 Layout** of the **Grouping** viewer to create an interactive analysis environment. In addition, this **Analysis Grid** viewer **Layout** displays by default if you are loading data from files in any of the following formats and the **File Sharing Perf SMB2/SMB Profile** is enabled in the **Options** dialog, as described in the Working With Message Analyzer Profiles topic:
    - *.etl
    - *.cap
    - *.pcapng
    - *.pcap

  See the **File Sharing Perf SMB2/SMB Profile** in the table of the indicated topic for more information about how to analyze the data in these files, where such analysis is enhanced by complementary **Layouts** for the **Grouping** and **Chart** viewers, which is described there.

  - **File Sharing SMB/SMB2** — provides a layout that groups by SMB **SessionId**, **TreeId**, and **FileName** to assist SMB troubleshooting.
  - **SMB Flat** — provides an SMB/SMB2 analysis environment that includes data columns such as **TimeDelta**, **SessionIdName**, **TreeIdNameReference**, **FileNameReference**, and **Header.MessageId**. After your data displays, click the **Flat Message List** button in the Filtering toolbar to remove Operations and simulate the Network Monitor view.
  - **SysLog** — provides an environment for analyzing SambaSysLogs with data columns such as **level**, **source_file**, **file_line**, **function**, and **content**. This **Layout** is intended to work with the **SysLog Layout** of the **Grouping** viewer to create an interactive analysis environment. In addition, this **Analysis Grid** viewer **Layout** displays by default if you are loading data from a SambaSysLog.log file and the **Samba Logs Profile** is enabled in the **Options** dialog, as described in the Working With Message Analyzer Profiles topic.

    See the **Samba Logs Profile** in the table of the indicated topic for more information about how to analyze the data in these logs, where such analysis is enhanced by complementary **Layouts** for the **Grouping** and **Chart** viewers, which is described there.

- **Network** category
  - **HTTP** — provides a **Layout** that is useful for troubleshooting HTTP issues, which includes the **ResponseTime**, **ContentType**, **StatusCode**, and **Summary** information columns. **ResponseTime** data is particularly useful because it can provide an indication of server performance in terms of the time it takes for the first server response to HTTP requests. **StatusCode** can quickly expose client and server errors and **ContentType** can provide an indication of the loads that the server is handling.
  - **Network Conversation Tree with Process ID** — provides a layout that groups by **DataSource**, **ProcessId**, **Network**, and **Transport**, to enable correlation of conversations with

process IDs in *.etl files.

- **Network Monitor** — displays the default Network Monitor view, with exception of several columns being named differently in Message Analyzer, but which are functionally equivalent to the corresponding columns in the default Network Monitor layout.
- **Process Name and Conversations** — provides a nested group configuration, where each top-level group node is identified by a **ProcessName** and nested groupings consist of **Network** and **Tranport** groups. Enables you to view the IP conversations and the TCP ports that carried those conversations for each process that is identified in the **ProcessName** group. This **Layout** also adds a **ProcessName** column to the **Analysis Grid** viewer column configuration.

  > **NOTE**
  > Message Analyzer can identify process names from .etl files that are generated with the Netsh utility. Therefore, you might use the **Process Name and Conversations** view **Layout** when you are working with event logs (*.etl), to expose data for the **ProcessName** group in the **Analysis Grid** viewer.

- **TCP Deep Packet Analysis with ABSOLUTE Sequence Number Flat** — provides a **Layout** that adds several columns to the default **Analysis Grid** viewer column configuration to expose the values of fields that can help you troubleshoot TCP and network layer issues. Added columns consist of the **TimeDelta**, **Flags**, **SourcePort**, **DestinationPort**, **PayloadLength**, **SequenceNumber**, **NextSequenceNumber**, **AcknowledgementNumber**, **WindowScaled**, **TopModule**, **Options**, and so on. Absolute sequence numbers are the long version of such numbers, as in the original format that is transmitted on the wire.
- **TCP Deep Packet Analysis with ABSOLUTE Sequence Number with Grouping** — provides a **Layout** that creates a hierarchy of **Network**, **Transport**, and **Sourceport** groups that isolate network conversations at top-level, the transport that carried them in the first nested group, and the ports over which the conversations transited in the second nested group. The **Network** and **Transport** groups provide quick access to data that can assist you in troubleshooting the network layer and related communication ports, for example IP addresses, conversation direction, and conversation ports. The columns in this **Layout** are identical to those of the **TCP Deep Packet Analysis with ABSOLUTE Sequence Number Flat**, as is the use of absolute sequence numbers.
- **TCP Deep Packet Analysis with RELATIVE Sequence Number Flat** — provides a **Layout** that exposes the values of fields that can help you troubleshoot TCP and network layer issues. The columns in this **Layout** add relative sequence number and relative block (**RelBlock**) columns to the ones previously described in the **TCP Deep Packet Analysis with ABSOLUTE Sequence Number Flat Layout**. Relative sequence numbers are the short version of such numbers for easier reading, as modified from the original wire format.
- **TCP Deep Packet Analysis with RELATIVE Sequence Number with Grouping** — provides a **Layout** that creates a hierarchy of **Network**, **Transport**, and **Sourceport** groups that isolate network conversations at top-level, the transport that carried them in the first nested group, and the ports over which the conversations transited in the second nested group. The **Network** and **Transport** groups provide quick access to data that can assist you in troubleshooting the network layer and related communication ports, for example IP addresses, conversation direction, and conversation ports. The columns in this **Layout** are identical to those of the **TCP Deep Packet Analysis with RELATIVE Sequence Number Flat Layout**, as is the use of relative sequence numbers.

> **NOTE**
> If you also specify a **TCP Viewpoint** with the TCP **Layouts** described here, you can isolate all TCP traffic as top-level messages for ease of analysis.

- **IIS** category
  - **IIS** — provides a **Layout** of data columns that is useful for troubleshooting Internet Information Server (IIS) logs. This **Layout** is intended to work with the **IIS Layout** of the **Grouping** viewer to create an interactive analysis environment. In addition, this **Analysis Grid** viewer **Layout** displays by default if you are loading data from an IIS.log file and the **IIS Logs Profile** is enabled in the **Options** dialog, as described in the Working With Message Analyzer Profiles topic.

    See the **IIS Logs Profile** in the table of the indicated topic for more information about how to analyze the data in these logs, where such analysis is enhanced by complementary **Layouts** for the **Grouping** and **Chart** viewers, which is described there.
- **Netlogon** category
  - **Netlogon** category — provides a **Layout** of data columns that is a basic environment for troubleshooting Netlogon logs. This **Layout** is intended to work with the **Netlogon Group by Message Type Layout** of the **Grouping** viewer to create an interactive analysis environment. In addition, this **Analysis Grid** viewer **Layout** displays by default if you are loading data from a Netlogon.log file and the **Netlogon Logs Profile** is enabled in the **Options** dialog, as described in the Working With Message Analyzer Profiles topic.

    See the **Netlogon Logs Profile** in the table of the indicated topic for more information about how to analyze the data in these logs, where such analysis is enhanced by complementary **Layouts** for the **Grouping** and **Chart** viewers, which is described there.
- **Common** category
  - **Perfmon Log** — provides a **Layout** with grouping for Perfmon logs that are saved as comma-separated value (CSV) files.
  - **Performance Top Down** — provides a **Layout** that is useful for analyzing performance issues associated with **TimeElapsed** and **ResponseTime** field values. In this **Layout**, the **TimeElapsed** data column is sorted in descending order (top-down) to expose Operations that are taking a long time to complete, for example with protocols that use request/response messages such as DNS, HTTP, and SMB. This may be an indication of a network issue if the corresponding server first response (**ResponseTime**) is a relatively low value. On the other hand, if the **ResponseTime** is a high value and the **TimeElapsed** is only slightly higher, this might indicate a slowly responding server rather than a network issue.
  - **Process View** — provides a **Layout** that enables top-down performance analysis, similar to the **Performance Top Down Layout**. However, this **Layout** also adds a **ProcessName** column for traces that contain process name information, so that you can correlate processes with **TimeElapsed** and **ResponseTime** data.
  - **ProcMon Logs** — provides a grouped **Layout** by process name and includes other data columns such as **Time Delta**, **PID**, **Operation**, and so on, for ProcMon logs.
  - **Protocol/Module Summary** — provides a grouped **Layout** by Module with nested message Types for a high-level overview of such data.
  - **WPP ETL** — provides a **Layout** with grouping and special data columns for WPP-generated events that are logged in a *.etl file.

- **My Column Layouts** category
  - **Raw Text Log** — provides a **Layout** that contains **MessageNumber** and **Summary** columns only. You can apply this **Layout** to the data that you load into Message Analyzer from any log file to obtain a high-level view of the log data. Note that none of the log fields that exist in the **Summary** column are provided a separate column for displaying their data in the **Analysis Grid** viewer.

    However, you may be able to parse the summary data into individual fields/columns, if a configuration file exists for the log type with which you are working. You can select different built-in configuration files from the **Text Log Configuration** drop-down list in the **New Session** dialog for a Data Retrieval Session after you add a log file to the files list. If it is a log type for which a configuration file already exists, the drop-down list is enabled for selection. For more information about working with text logs, see Opening Text Log Files.
- **My Items** category
  - **Example View Layout** — provides a sample **Layout** based on TCP fields that you can modify as you wish. You can save your changes with a new **Layout** name that is appropriate for the type of analysis it supports.

Managing View Layouts

The **Layout** asset collection Library drop-down list also provides a set of commands that enable you to manage your view **Layouts**. Most commands are accessible in a submenu that displays when you click the **Manage Layouts** item in the **Layouts** drop-down list, as described in the "Manage Layout Commands" section below. However, the following frequently used command is directly accessible from the **Layout** drop down list for convenience:

- **Save Current Layout As...** — enables you to save view **Layouts** that reflect unique column configurations that you create with **Field Chooser**, so that they appear as items in a specified **Category** in your local asset collection **Layout** Library. Thereafter, you can reapply them whenever you want to by selecting them from the **Layout** Library, or you can export them to a designated location for sharing with other users.

  When you use this command to save the current **Layout**, the **Edit Item** dialog displays to enable you to specify a **Name**, **Description**, and **Category** in which to place the **Layout**. Note that any **Layouts** that you save from the **Edit Item** dialog will display in a subcategory under the top-level **My Items** category.

Manage Layout Commands

Other commands that are available for managing view **Layouts** display in a submenu that appears when you select the **Manage Layouts** item in the **Layout** drop-down list, as follows:

- **Save Current as Default User Layout** — enables you to save the current **Analysis Grid** viewer column configuration as the default **Layout** for the **Analysis Grid**, which then displays whenever you specify the **Analysis Grid** as the viewer for session results, providing that an enabled **Profile** does not override that **Layout**.
- **Load Default User Layout** — enables you to restore the **Layout** that you saved as the default with the **Save Current as Default User Layout** command, as needed.
- **Restore Application Default Layout** — enables you to restore the default view **Layout** for the **Analysis Grid** that ships with Message Analyzer, as described in the Analysis Grid Viewer topic.
- **Manage...** — displays the **Manage Column Layouts** dialog from where you can export one or more **Layouts** as shareable items so that other users can take advantage of your **Layout** configurations. You

can also import one or more **Layouts** that other users have made available as shareable items, to expand your local **Layout** Library. In addition, you can edit the information for any column **Layout** in the **My Items** category of the dialog by executing the **Edit** command from the dialog's right-click context menu.

Sharing View Layout Items

When you save a view **Layout**, it becomes part of a local asset collection Library containing **Layout** items that you can manage and share with others. Message Analyzer provides a simple way to expose these **Layout** items to others for sharing, or to retrieve **Layouts** that others have shared. From the **Manage Column Layout** dialog, you can share **Layout** items directly with others by selecting the item/s you want to share and then clicking the **Export** button on the dialog tool bar. The **Save Library** dialog then displays so that you can enter **Title**, **Description**, and **Organization** data. Thereafter, you can specify the file share location where you want to post your items. You can also use the **Import** feature in the same dialog to access **Layout** items that have been shared by others to a designated file share. The **Manage Column Layout** dialog is accessible by selecting the **Manage...** item in the submenu that appears when you select the **Manage Layouts** item from the **Layout** asset collection Library drop-down list on the **Analysis Grid** viewer toolbar.

You can also share your **Layout** items through the Message Analyzer Sharing Infrastructure by creating a user feed from the **Settings** tab of the **Asset Manager** dialog, which is accessible from the Message Analyzer global **Tools** menu. When you create your own user feed, you will point it to a file share or other designated location where you post your **Layout** items. To post your items, click the **Export** button on the **Manage Column Layout** dialog tool bar and navigate to the designated share location. Thereafter, you can update existing items or add others and make them available to team members or other users through the configured feed, where users can view, download, and synchronize to updates to your **Layout** items. However, the synchronization feature will be available only in a future Message Analyzer release, at which time the Sharing Infrastructure publishing features will enable others to automatically synchronize to any updates that you make to your **Layout** items on a user feed. In the meantime, you can perform a manual configuration process that enables users to synchronize to your **Layout** asset collection updates, as described in Manual Item Update Synchronization.

Microsoft also provides a default **Message Analyzer** feed on the **Downloads** tab of the **Asset Manager** dialog that enables you to download **Message Analyzer View Layouts** or **Azure Storage View Layouts** asset collections from a Microsoft web service and to synchronize with asset collection updates that are periodically pushed out by the service. At any time, you can perform a download of an auto-synced collection from the **Settings** tab of the **Asset Manager** dialog.

> **NOTE**
>
> **Message Analyzer View Layouts** that apply to the **Analysis Grid** viewer and **Message Analyzer Grouping View Layouts** that apply to the **Grouping** viewer are separate and function independently of each other. You can view these asset collections in the **Asset Manager** dialog.

More Information

To learn more about sharing asset collections in Message Analyzer, including further details about the common **Manage <*AssetType*>** dialog, see the Sharing Infrastructure topic.
To learn more about the **Field Chooser**, see the Field Chooser Tool Window topic.
To learn more about the **Analysis Grid** viewer **Group** function, see Using the Analysis Grid Group Feature.
To review some simple examples of TCP troubleshooting with Message Analyzer, see Procedures: Using the Data Filtering Features.

Using the Analysis Grid Group Feature

8 minutes to read

Because there are many different message conversations that can take place at different network layers, obtaining a view of your trace data that is relevant to resolving a particular diagnostic issue can be challenging at times. Also, since protocol analyzers often handle large amounts of data where messages of interest can be scattered, it can be difficult and time-consuming to find specific data that you want to examine. To accommodate these challenges, Message Analyzer provides a data **Group** function in the **Analysis Grid** viewer that enables you to bubble up and organize relevant and important data into a grouped display. You can access and apply the data **Group** function as described in Grouping Operations later in this section.

Reorganizing Data with the Group Function

The **Group** function enables you to create a high-level view of your data that is arranged by groups, to provide an efficient way of presenting specific information from high-volume traces. This can include grouping the data values of a particular entity such as a property, field, or annotation displaying as an **Analysis Grid** column or in the **Details Tool Window**. You can also create nested groups of multiple data values. In Message Analyzer, a group is an expandable and collapsible node that is labeled and contains the count of either child nodes or a set of child messages that are grouped by a data field or property value that is defined in OPN. Message Analyzer enables you to customize and accelerate your data analysis process by the strategic application of multiple grouping operations that display in a hierarchical manner. By carefully selecting which data is to be grouped in a set of groups, and by considering the data on which to pivot (the order in which you apply grouping operations), you can quickly isolate specific traffic and simplify the analysis process.

For example, in a particular set of trace results, you might add the **ProcessId** (ETW), **Network**, and **Transport** columns to the **Analysis Grid** viewer with the **Field Chooser Tool Window** and then perform **Group** commands on each column in the indicated order. The result would look like the following:

- **ProcessId Groups** — consist of separate Groups of expandable nodes that each consolidate identical ETW ProcessIds from various components whose events were captured by Message Analyzer. Each **EtwProviderMsg.EventRecord.Header.ProcessId** node indicates the **ProcessId** value and specifies the number of child/nested **Network** nodes (groups) in parentheses. You will find this field under the **ETW** node in **Field Chooser**.
- **Network Groups** — consist of separate nested Groups of expandable nodes that each consolidate identical pairs of **Source** and **Destination** addresses for IP conversations. Each **Network** node indicates the IP **Source** and **Destination** address and specifies the number of child/nested **Transport** nodes (groups) in parentheses. You will find this field under the **IPv4** node in **Field Chooser**.
- **Transport Groups** — consist of separate nested Groups of expandable nodes that each consolidate identical TCP or UDP **SourcePort** and **DestinationPort** pairs where the IP conversations took place. Each **Transport** node indicates the associated TCP or UDP port values and specifies the number of messages that are contained in the node. By expanding the **Transport** node, you can examine the messages that met the filtering criteria of the grouping hierarchy that you created. You will find this field under the **TCP** node in **Field Chooser**.

After you complete the **Group** commands on these columns, you should see a display similar to the following when various Group nodes are expanded:

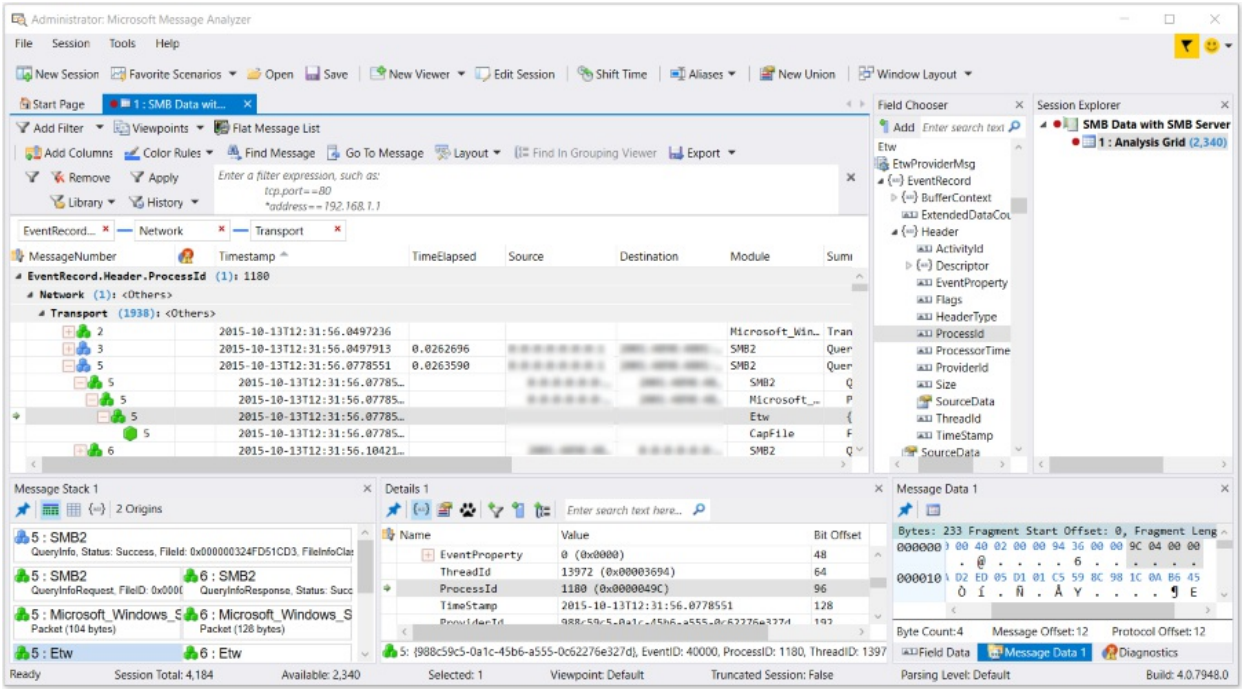

Figure 42: Message Analyzer Analysis Grid grouping operation

As shown in the preceding figure, each group is designated by a label that contains the name of the column to which you applied the **Group** command. Note that only the last group that you create contains messages that have been filtered down by the criteria that created each parent group. The displayed groups isolate messages based on TCP port data and IP addresses under ProcessIds, so that you can quickly scan the IP conversations that took place across source and destination node ports for the particular events that Message Analyzer captured.

Another example of grouping consists of applying a **Group** command to the **DiagnosisTypes** column in the **Analysis Grid** viewer to bubble up errors that occurred in a trace. The Groups that display are each designated by a specific diagnostic type such as **Application**, **Insufficient Data**, **Parsing**, or **Validation**. Note that you can also examine a summary of diagnostic errors that occur in a trace by opening the **Diagnostics Tool Window** from the **Windows** submenu of the global Message Analyzer **Tools** menu. If the **Diagnostics** window does not appear in this menu, you will need to enable it on the **Features** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. After you select the **Diagnostics** feature and click **OK** to exit the **Options** dialog, a Message Analyzer restart is required. Note that the **Diagnostics** window is currently preview feature.

> **NOTE**
>
> You can also add a view **Filter** or sort columns while grouped data is displayed in the **Analysis Grid** viewer. For example, you could add a **Filter** such as the following to the text box of a Filter panel on the Filtering toolbar to remove all message groups except those that contain the specified IP address: `IPv4.address==192.168.1.1`. You can even apply the **Group** command while Message Analyzer is capturing messages.

Grouping Operations

You can group data from **Analysis Grid** viewer columns by selecting the **Group** or **Group By Multiple Values** menu commands that display when you right-click an **Analysis Grid** viewer column. You can also group data from the **Details** window below the **Analysis Grid** viewer by selecting the **Add 'fieldName' as Grouping** context menu command that displays when you right-click a row in the **Details** window. The *fieldname* value in this command is a placeholder for an actual field name in the **Name** column of **Details**.

Using the Analysis Grid Group Command

The **Group** command creates separate groups based on varying **Column** values in the **Analysis Grid** viewer or field values in the **Details** window. When you apply the **Group** command to a particular **Column**, each differing value under that **Column** for each top-level message row is consolidated into a separate group of identical values, while the message origins tree associated with each top-level message or operation is preserved for the sake of maintaining context. For example, if you right-click the **Module** column and select the **Group** command, Message Analyzer creates a separate Group for each protocol module and populates that Group with all the top-level messages or operations that have the identical **Module** name. Message Analyzer identifies the results of the grouping operation with a label that designates the name of the column to which you applied the **Group** command.

Because Message Analyzer supports the application of **Group** commands to multiple columns, you can perform grouping operations on as many data columns as you want, to achieve different perspectives of your data. Each time you apply a new **Group** command to a different data column, the results of the previous grouping operation are modified to include the new grouping operation as a nested **Group** of filtered messages within the existing parent **Group**. When this occurs, the criteria for the new grouping operation are based on the values of the existing **Group**. For example, if you performed one grouping operation based on the **Source** column and then did another grouping operation based on the **Destination** column, the results will display all identical **Destination** addresses for each **Source** address Group in a separate Group that is nested in the parent **Source** address Group. The next grouping operation applies the same type of logic when you perform another **Group** command on a different column.

Using the Group By Multiple Values Command

The **Group By Multiple Values** command creates groups based on varying values in a particular field that might be different at various stack levels. For example, you might see IP addresses at the top-level in the **Source** column of the **Analysis Grid** viewer, while in the stack there are also hidden Ethernet addresses. By using the **Group By Multiple Values** command on the **Source** column, you can organize the IP and Ethernet address values into separate groups for enhanced data analysis capabilities.

Exposing and Hiding Grouped Data

To expose the data contained in any particular Group, you can click the arrow node to the left of each Group. To expose the data contained in all Groups, you can right-click the Group label and select the **Expand All Groups** menu command. To hide the data contained in all Groups, you can select the **Collapse All Groups** command from the same menu.

Reorganizing Grouped Data

If you create more than one data Group in the **Analysis Grid** viewer, you have the option to reorganize the Groups and subsequently rearrange the nesting order of the nodes that contain your data. By reorganizing the Groups, you can achieve a different analysis perspective on the data. To accomplish this, you can drag and drop any Group label to the left or right of any other Group label, as appropriate, and the data grouping will be reorganized as if you had originally performed the **Group** operation in that order. For example, if you **Group** the **Source** column in the **Analysis Grid** viewer first and then **Group** the **Destination** column, the **Destination** addresses will be contained in **Destination** nodes that are nested under the **Source** nodes. However, if you drag and drop the **Source** group label to the right of the **Destination** group label, the grouping will be reversed with the **Source** addresses contained in **Source** nodes now nested under the **Destination** nodes.

Removing Grouped Data

To undo a Grouped display configuration, click the **x** mark in the Group label that identifies the grouped column data you want to remove. As you remove groups in the indicated manner, data for the selected Group is removed from the display and any remaining Groups are filtered and reorganized according to the hierarchy established by the original grouping operation.

See Also

Grouping Viewer

Viewing OPN Source Code

2 minutes to read

Message Analyzer provides a definition viewer that displays Open Protocol Notation (OPN) description code for message modules that display in the **Analysis Grid** viewer and for most fields that display in the **Details Tool Window**. After you load data from a native Message Analyzer trace file or after you stop a Live Trace Session, you can view the OPN definition for most message fields, although there are exceptions such as CSV/TSV and some text-based log files for which there might be no OPN definitions. This feature is primarily intended to help developers or other interested users to understand more about OPN descriptions and their use as message parsers.

Working with OPN Definition Data

You can display the OPN definition viewer as a separate tab in an Analysis Session by right-clicking a message in the **Analysis Grid** viewer, a field in the **Details** window, or a field in **Field Chooser Tool Window**, and thereafter selecting the **Go To 'fieldName' Definition** command to show the OPN viewer and the OPN definition for the associated field or module. Only when you open the OPN definition viewer from the **Analysis Grid** or the **Details** window is an *entityName* designation included in single quotation marks in the **Go To Definition** command; otherwise it is absent. This designation corresponds to the module or field name that is associated with the message row or field row that you right-click, respectively. The entityName will therefore be identified by a message type in the **Module** column of the **Analysis Grid** viewer or by a field **Name** in the **Details** window. After you perform this action, the OPN definition viewer displays the read-only OPN code that defines the associated module or message field you selected.

Within the OPN definition viewer, there are several commands that you can access by right-clicking any location inside the viewer window. These commands and the actions they perform are described as follows:

- **Find** — displays a search control bar that enables you to enter text characters that you want to search for in the OPN code.

  As soon as the OPN viewer is open, a search control bar displays with the selected field information. You can then use the forward and backward arrow keys to search through the OPN code in the direction you specify for occurrences of the search characters. The search control bar also provides check box options for **Match case**, **Match whole words**, and **Use regular expressions** in a drop-down, which establishes the indicated criteria for the search. If you want to search for another entity in the displayed definition, simply type the entity name while the default search characters are highlighted and then use the search arrow keys in the control to locate any occurrences of the entity.

- **Search on Web** — enables you to search on the web for a highlighted string of characters that you selected.

- **Search in Microsoft Protocol Documentation** — enables you to search for protocol-related information in the Windows Protocols documentation in the MSDN Library or other locations on the web.

  For example, you might search for the specification of a particular protocol such as AIPS.

More Information

To learn more about OPN, see the OPN Programming Guide.

Viewing Session Statistics and Progress

6 minutes to read

Message Analyzer provides a robust set of monitoring facilities such as progress indicators and various types of session statistics that include hover-over tool tips. For example, Message Analyzer enables you to view the progress of various message retrieval operations such as the following:

- Data loaded through the **Open** feature on the Message Analyzer **Start Page**.
- Data loaded from the **Recent Files** feature from the Message Analyzer global **File** menu.
- Data loaded through a Data Retrieval Session.
- Data captured in a Live Trace Session.

Message Analyzer also enables you to view progress indications for multiple concurrent sessions that load data at the same time. You can also observe the progress of message processing during data manipulation operations, which includes the application and removal of view **Filters**, **Groups**, **Viewpoints**, **Pattern Matches**, **Layouts**, and **Time Filters**. In addition, you can hover over any top-level session node in **Session Explorer** with your mouse and view the message count for the associated session in a tooltip. You can also hover over any session viewer node in **Session Explorer Tool Window** and display a tooltip that indicates the analytical assets that are applied to that viewer.

Indicating Progress and Statistics

Visual progress indications and statistics for the previously described Message Analyzer operations are displayed in the **Session Explorer** window. To provide these indications when loading data, capturing data, or manipulating data, **Session Explorer** utilizes the following progress and statistics indicator components:

- **Progress bar control** — provides the following types of progress indications:
  - **Relative** — displays a continuous solid bar that propagates from left to right when the number of messages in a message source is known, for example, when loading data from .cap, .pcap, .matu, and .log files.
  - **Ambiguous** — displays a short block that scrolls across the control in marquee fashion when the number of messages cannot be determined, for example, when capturing data live or when loading data from .etl, .matp, .evtx, .tsv, and .csv files.
- **Relative progress label** — to the right of each progress bar control in a **Session**, a parenthetical label can provide numerical percent-complete statistics for an operation. Percent-complete values display only when the total number of messages in a message source is known, which is therefore considered a *relative* type progress indication.
- **Session node label** — to the right of each **Session** node in **Session Explorer**, there is a parenthetical label that is populated with the total number of messages processed and available to the current **Session**. If you hover your mouse over a **Session** node, the following tool tip is displayed:
  "The value represents the number of messages currently available for this session".
- **Viewer node label** — to the right of the **Session Explorer** node that contains the data viewer that you specified when starting a Data Retrieval Session or Live Trace Session, there is a parenthetical label that indicates the number of parsed messages that are available to the session viewer that you specified, for example, the **Analysis Grid** or **Gantt** viewer. Note that when you load data through the **Open** feature, the

Analysis Grid viewer displays by default, unless you have changed the default viewer in the **Default Viewer** pane on the **Profiles** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. Also, if you hover your mouse over a viewer node, the following tool tip is displayed: *"The value represents the number of messages currently available for this view"*.

- **Hover-over tool tips** — as indicated earlier, tooltips that display when you mouse-hover over any top-level session node or any session viewer sub-node in **Session Explorer**, provide a quick indication of session message count and the analytical assets that you have applied to a viewer, respectively. Applied assets can include a **Viewpoint**, **Viewpoint Filter**, view **Filter**, and/or a **Message Range Filter**.

> **NOTE**
> Session viewer tabs also display the same tool tip information when you hover over them with your mouse.

- **Session color coding** — each top-level session node and related session viewer subnodes in **Session Explorer** are color-coded for quick identification. Different sessions are identified with a different colored dot, which helps to automatically organize your session data viewers.

> **TIP**
> Session viewer tabs also contain associated colored dots, to enable you to quickly correlate all viewers of the same session.

> **NOTE**
> **Session Explorer** also contains different **Session** node icons that indicate the type of session to which progress statistics apply, for example, a Live Trace Session versus a Data Retrieval Session. In addition, icons are added to the right side of session viewer subnodes to indicate the application of an analytical asset. For example, a funnel shaped glyph indicates when a **Session Filter** or view **Filter** was applied to the messages in your Live Trace Session or Data Retrieval Session.

Observing Progress Indicator Behaviors

The behavior of **Session Explorer** progress indicators varies depending on the tasks you are performing, as follows:

Loading parsed messages — when you load data from a .matp file in a Data Retrieval Session, the progress bar might display an *ambiguous* progress indication for larger files, by scrolling a block across the control in marquee fashion, along with the following simultaneous indications:

- **Session** node indicator above the progress bar control — shows the total number of messages processed in the session.
- **Viewer** indicator below the progress bar control — shows the number of parsed messages available in the specified data viewer for further processing.

> **NOTE**
> Because a .matp file is already parsed, it should load very fast, and in this case a progress indication might not display.

Loading unparsed messages — when you load data from a .matu file in a Data Retrieval Session or through the **Open** and **Recent Files** features on the **Start Page**, the progress bar control displays a continuous solid bar that propagates from left to right, while providing the identical indications described in the previous list, consisting of the total number of messages processed (above the control) and the number of parsed messages available in the viewer (below the control). It also provides progress as a percent-complete value, which indicates the number of

messages processed/parsed out of the total message count in the .matu file.

Capturing data live — when you are capturing messages in a Live Trace Session, the progress bar control displays a scrolling block across the control (until such time that you stop the Live Trace Session) along with simultaneous indications of the number of captured messages (above the control) and the total number of parsed messages that are available in the specified viewer (below the control).

> **NOTE**
>
> **Session Total** and **Available** message indicators are also provided in labels on the Status Bar of the Message Analyzer UI and reflect the identical values that display in the **Session** node and viewer node labels, respectively.

Manipulating data

Progress indications are also displayed whenever you apply or remove view **Filters**, **Groups**, **Viewpoints**, **Pattern Matching**, view **Layouts**, and **Time Filters** to or from a set of displayed messages, respectively, or when you toggle **Operations** from the **Viewpoints** drop-down list on the Message Analyzer Filtering toolbar. Progress indications are also provided when you sort columns. The progress indications that display for each of these data manipulation operations consists of:

- A continuous solid progress bar that propagates from left to right.
- A corresponding relative progress indication that displays a percent-complete value, along with an operation indicator such as *Filtering*, *Sorting*, *Applying Viewpoint*, and so on.
- The number of messages processed and available to the specified viewer based on the applied data manipulation criteria.

> **NOTE**
>
> The **Diagnostics Tool Window** also contains its own progress indicator; however, it usually activates only when a large set of trace results is being scanned and diagnosis messages are being aggregated into summaries that thereafter will display in this tool.

You can also apply simultaneous data manipulation operations across multiple sessions and observe progress indications that apply to each session. However, if you apply multiple operations in the same session viewer, the order in which operation indicators display is determined by a preset order of precedence that Message Analyzer applies, regardless of the sequence in which you start data manipulation operations.

See Also

Session Explorer Tool Window

Grouping Viewer

45 minutes to read

To augment your analysis capabilities, the **Grouping** viewer enables you to organize your traffic into summary hierarchies based on Grouping viewer **Layouts** that contain predefined message field Groups that exist in nested configurations. The **Grouping** viewer enables you to extract specific types of data from any large data set and organizes this information into a hierarchy of one of more nested Groups that provide instant access to data that you can interactively correlate with **Analysis Grid** viewer message details.

The **Grouping** viewer is accessible from the **New Viewer** drop-down list, which in turn is accessible from the following locations:

- Global Message Analyzer **Session** menu.
- Global Message Analyzer toolbar.
- **Session Explorer** right-click context menu.

The **Layout** that displays in the **Grouping** viewer by default has a grouped configuration that consists of **Network** and **Transport** groups, although many other layouts are available such as the **File Sharing SMB/SMB2**, **TCP Deep Packet Analysis**, and **Process Name and Conversations Layouts**, the latter of which organizes data in a way that is similar to the Network Monitor **Conversation Tree**. Note that many **Grouping** viewer **Layouts** are designed to work with specific **Analysis Grid** viewer **Layouts** to create an interactive analysis environment through Message Analyzer **Profile** configurations that apply such **Layouts** to these viewers, as described in Working With Message Analyzer Profiles.

What You Will Learn

In this section, you will learn about the functions and features of the **Grouping** viewer, as described in the following topics:

Utilizing the Grouping Viewer Capabilities
Changing the Analysis Perspective Through Group Reorganization
Grouping Viewer Layouts
Understanding the Built-In Grouping View Layouts
Manipulating Group Displays
Grouping Viewer Modes of Operation
Locating Analysis Grid Messages in the Grouping Viewer
Grouping Viewer Display Features
Adding New Groups
Editing a Built-In Layout
Creating a Grouping Viewer Layout Template
Managing Grouping View Layouts

Utilizing the Grouping Viewer Capabilities

An important capability that is available for the **Grouping** viewer is that you can create custom grouped layout configurations of your own based on message field groups that you choose and group nesting configurations that you create, to obtain the analysis perspectives that are most useful in your environment. In addition, by manually altering the way your message field Groups are nested, you can adjust (pivot) your grouped layout on-the-fly to acquire different message correlation configurations that result in unique analysis contexts. A major advantage of the **Grouping** viewer, is that you can organize data into unique

hierarchies to expose targeted information that you can quickly extract from a large data set, which otherwise would be difficult to achieve. More specific advantages include the ability to do the following:

- **Locate the Group(s) that contain the largest traffic volumes**:

  For example, in the **Process Name and Conversations Layout** of the **Grouping** viewer, you can compare the traffic volumes across all the top-level **ProcessName** Groups to determine the Groups with the most traffic.

- **Isolate all messages to a top-level Group type and drill down further for data in nested Groups, to obtain a concise analytical focus on specific messages of interest**:

  In the **Process Name and Conversations** layout, this means you can isolate messages from one or more process names, view the distribution of process ID traffic for each process name, and drill down further to view the related **Network** layer conversations and the **Transport** layer ports that carried them.

- **Correlate messages in the Analysis Grid viewer with the Groups in which they appear in the Grouping viewer**:

  For example, you can use the **Find in Grouping Viewer** command that is accessible as a right-click context menu item for any message in the grid, to locate the corresponding Group in which the selected message appears.

- **Correlate messages across different data sources**:

  For example, you might compare messages from a combination of associated log files and live trace results.

Changing the Analysis Perspective Through Group Reorganization

The **Grouping** viewer enables you to reorganize how the Groups are nested in order to achieve a different analysis perspective. For example, consider the Group hierarchy of the **Process Name and Conversations** layout shown in the upper left sector of the user interface in the figure that follows. The particular groupings shown in this figure isolate traffic by **ProcessName**, **ProcessID**, **Network** layer, and **Transport** layer, as previously indicated. This Group configuration enables you to correlate the TCP/UDP ports that carried IPv4/IPv6 conversations across various process IDs for corresponding process names. To achieve a different perspective on the same data, you can quickly pivot the data by dragging and dropping any Group label into a different position in the **Grouping** viewer toolbar area to automatically reorganize the data according to a new Group nesting configuration. By doing this, you can view the data in a different way and potentially expose issues that were formerly less apparent.

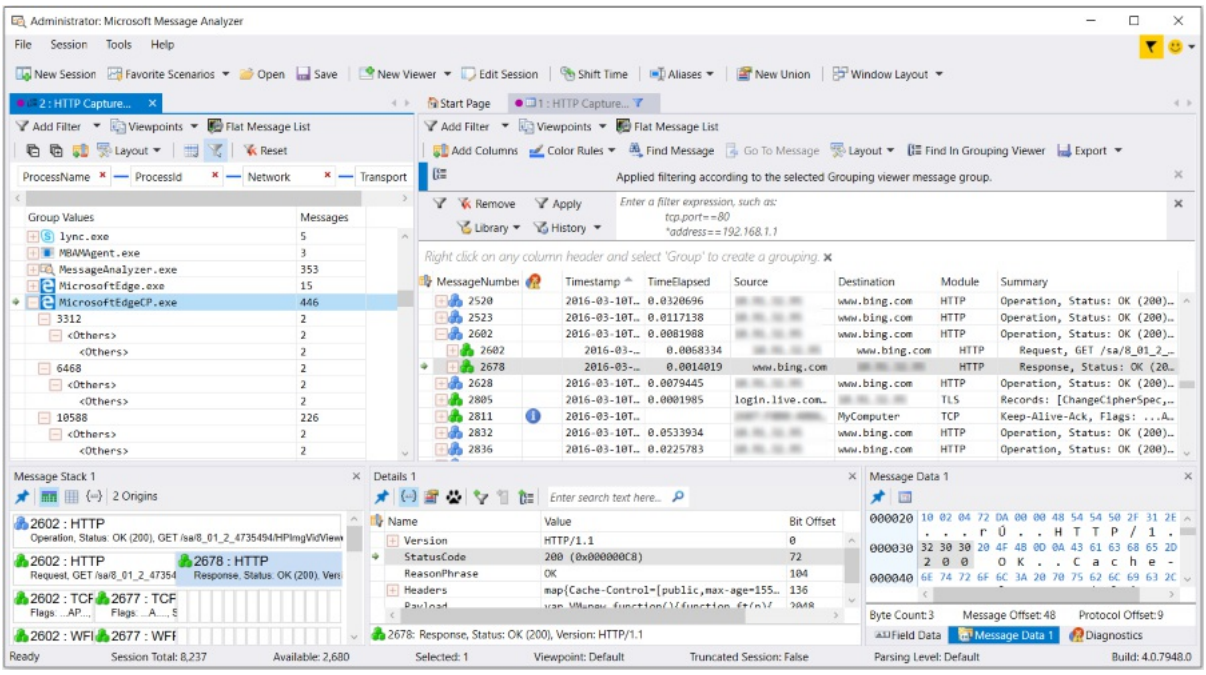

Figure 42: Message Analyzer Grouping Viewer with Process Name, ProcessId, and Conversations Layout

---

More Information

To learn more about the **Process Name and Conversations** layout for the **Grouping** viewer, see Understanding the Built-In Grouping View Layouts.

To learn more about reorganizing Grouped information, see Manipulating Group Displays.

Grouping Viewer Layouts

The **Grouping** viewer toolbar contains a **Layout** drop-down list that enables you to select different built-in **Layouts** that you can use for data analysis purposes. Note that the assets you access from this **Layout** drop-down list apply to the **Grouping** viewer only and are available only when the **Grouping** viewer is displayed.

You can modify any of the built-in **Layouts** and save your changes if you want, and you can set any **Layout** as the default for the Grouping viewer. You also have the option to create your own layout with the use of the **Field Chooser Tool Window** and save it as indicated as an item in the **Message Analyzer Grouping View Layouts** asset collection Library, which is exposed by the **Layout** drop-down. Briefly, you can create a new Grouping **View Layout** in either of the following ways:

- Modify a currently displayed **Layout** by adding or removing Groups and then select the **Save Current Layout As...** command from the **Layout** drop-down list on the **Grouping** viewer toolbar to display the **Edit Item** dialog, from where you can specify **Name**, **Description**, and **Category** information. When you click **OK** to exit the dialog, Message Analyzer saves your custom **Layout** in the **Category** name that you specified, under the top-level **My Items** category of the **Layout** Library.

> **NOTE**
>
> You can add Groups by using any of the methods that are described in Adding New Groups. You can remove Groups by simply clicking the **x** mark in one or more Group labels, just below the **Grouping** viewer toolbar.

- Right-click an item in the **Layout** drop-down list and then select the **Create a Copy** command to display the **Edit Item** dialog, where you can modify **Name**, **Description**, and **Category** information. After you click **OK** to exit the **Edit Item** dialog, Message Analyzer saves your custom layout in the **Category** name

that you specified, under the top-level **My Items** category of the **Layout** Library. Thereafter, you will have access to several context menu commands that display whenever you right-click a **Layout** in the **My Items** category, where all user-created **Layouts** for the **Grouping** viewer are saved by default. Thereafter, you can display and customize your **Layout** by adding or removing Groups and saving your changes as previously described.

> **NOTE**
>
> To support the requirements of your environment, you can create and save as many different **Layouts** for the **Grouping** viewer as you wish. To learn more about creating Grouping **Layouts**, see Creating a Grouping Viewer Layout Template.

Understanding the Built-In Grouping View Layouts

The **Message Analyzer Grouping View Layouts** asset collection Library is accessible from the view **Layout** drop-down list on the **Grouping** viewer toolbar. It contains several built-in **Layouts** that enable you to group and summarize messages based on field values that are specifically exposed by various **Layouts**, to create alternate analysis perspectives on different types of data. It also enables you to expose data that is deeply hidden in large volume traces through the filtering action of nested group configurations. The built-in **Layouts** are described in this section.

> **NOTE**
>
> The **Layout** drop-down list on the **Grouping** toolbar pertains to the **Grouping** viewer only. It is not associated with the **Layout** feature that is accessible from the **Analysis Grid** viewer toolbar. On the other hand, many of the **Grouping** viewer **Layouts** that are described in this section are specifically intended to work with **Analysis Grid** viewer **Layouts** to create a unified interactive analysis environment, as mentioned earlier.
>
> You can facilitate interaction between these viewers by clicking Group nodes to display corresponding messages in the **Analysis Grid** viewer for further scrutiny of message details. How these messages display depends on the Grouping viewer mode, as described in Grouping Viewer Modes of Operation. Another interaction consists of locating an **Analysis Grid** message within the Group hierarchy in the **Grouping** viewer, as described in Locating Analysis Grid Messages in the Grouping Viewer.

- **Cluster** category
  - **Cluster Logs** — this grouping **Layout** provides a top-level **InfoLevel** Group that contains all the debug levels generated by the Cluster Service and any of its subcomponents. Debug levels consist of DBG, ERR, INFO, and WARN. **Subcomponent** is the first nested Group under **InfoLevel** and it is populated with acronyms that represent various subcomponents of the Cluster Service such as the Global Update Manager (GUM), Failover Manager (FM), and Database Manager (DM). Nested below the **Subcomponent** Group is the **ProcessId** Group, which exposes the hexadecimal identifier of the executing processes launched by a particular Cluster Service component. With these three Groups, you can identify the components in which specific errors and warnings occurred during Clustering operations, along with the IDs of associated executing processes. You can then select any of these Groups to expose the log entries associated with each Group in the **Analysis Grid** viewer for further analysis.

    This grouping **Layout** is intended to work with the **Cluster Log Layout** of the **Analysis Grid** viewer to create an interactive analysis environment. You will be able to correlate the data most effectively if you have this **Analysis Grid** viewer **Layout** displayed. For example, if you drill down to the **ProcessId** Group and select process IDs, you can display the associated log entries in the **Analysis Grid** viewer for analysis of log entry details. With the **Cluster Log Layout** in the **Analysis Grid**, you will also have access to other log fields of interest such as

RemainingText, which can expose the cause of errors and warnings, or provide other debugging information.

Note that if you have the **Cluster Logs Profile** enabled on the **Profiles** tab of the **Options** dialog, the **Cluster Log Layout** automatically displays in the **Analysis Grid** viewer after you load data into Message Analyzer from a Cluster.log file, as described in the Working With Message Analyzer Profiles topic. However, if you want to display the **Cluster Logs Layout** in the **Grouping** viewer, you will need to manually select the **Default** item in the **Grouping** drop-down list that is accessible from the **New Viewer** drop-down list. Otherwise, if the **Cluster Logs Profile** is not enabled, you will need to manually select the **Cluster Logs Layout** in the **Grouping** drop-down list, rather than the **Default** item, to use this **Layout**. The same is true of the **Chart** viewer **Layout** that is also part of the **Cluster Logs Profile**.

For more information about analyzing Cluster logs with Message Analyzer, where such analysis is enhanced by complementary **Layouts** for the **Analysis Grid** and **Chart** viewers, see the **Cluster Logs Profile** in the table of the previously indicated topic.

- **Common** category
  - **ETW Guids and IDs** — this grouping **Layout** provides a top-level **ProviderId** Group with a nested **Descriptor.Id** Group that enables you to obtain a quick assessment of the event volumes associated with each ETW provider that participated in the trace, along with IDs of the events that each provider wrote to an ETW session. The **ProviderId** data specifies the GUID of the ETW trace provider that generated an Event and the **Descriptor.Id** data specifies Event identifiers, which are part of an Event Descriptor, as described in the ETW Framework Conceptual Tutorial topic. You can isolate the events per provider or individual event IDs by clicking a group of interest.

    This grouping **Layout** is intended to work with the **ETW Layout** of the **Analysis Grid** viewer to create an interactive analysis environment. You will be able to correlate the data most effectively if you have this **Analysis Grid** viewer **Layout** displayed. For example, if you select any item in the **Descriptor.Id** Group under a **ProviderId** Group, you will be able to see additional details in the **Analysis Grid** viewer for one or more associated log entries. You will also be able to review data in the **Summary** column of the **Analysis Grid** viewer for additional information such as error and failure descriptions.

    Note that if you have the **ETW Analysis Profile** enabled on the **Profiles** tab of the **Options** dialog, the **ETW Layout** automatically displays in the **Analysis Grid** viewer after you load data into Message Analyzer from a *.etl file, as described in the Working With Message Analyzer Profiles topic. However, if you want to display the **ETW Guids and IDs Layout** for the **Grouping** viewer, you will need to manually select the **Default** item in the **Grouping** drop-down list that is accessible from the **New Viewer** drop-down list. Otherwise, if the **ETW Analysis Profile** is not enabled, you will need to manually select the **ETW Guids and IDs Layout** in the **Grouping** drop-down list, rather than the **Default** item, to use this **Layout**. The same is true of the **Chart** viewer **Layout** that is also part of the **ETW Analysis Profile**.

    For more information about analyzing event trace logs with Message Analyzer, where such analysis is enhanced by complementary **Layouts** for the **Analysis Grid** and **Chart** viewers, see the **ETW Analysis Profile** in the table of the previously indicated topic.

  - **Event Viewer** — this grouping **Layout** provides a top-level **ProviderName** Group with successively nested **Level**, **Channel**, and **EventID** Groups, that enable you to obtain a quick assessment of the event volumes associated with each ETW provider that participated in the trace, any error levels that occurred for the events issued by such a provider, and the IDs of those events. By selecting any of these Groups that are described in the list that follows, you can isolate the corresponding event log entries in the **Analysis Grid** viewer, to further

scrutinize the details. Note that you will find data for each of these fields in the **Details Tool Window** after you select any particular log entry in the **Analysis Grid** viewer following Group selection in the **Grouping** viewer:

- **ProviderName** — this top-level field is the name of the ETW provider that raised events and wrote them to the ETW session from which your data is displaying.
- **Level** — this field can include error Levels in the range of 1-5, for example, Critical (1), Error (2), Warning (3), Information (4), and Verbose (5).
- **Channel** — this field displays the target audience for the event/s, which is typically specified in an ETW provider manifest.
- **EventID** — this field specifies the ID for events that were written by an ETW provider.

This grouping **Layout** is intended to work with the **Event Log Layout** of the **Analysis Grid** viewer to create an interactive analysis environment. You will be able to correlate the data most effectively if you have this **Analysis Grid** viewer **Layout** displayed. For example, this **Analysis Grid** viewer **Layout** enables you to perform deep analysis of ETW events with the **ProcessId**, **ThreadId**, **Level**, **ActivityId**, **Channel**, **Version**, **OpCodeDisplayName**, and **EventData** columns of the **Layout**, in addition to the fields described in the previous list. By selecting any particular Group, you can immediately expose and correlate all this data, including **Details**, for specific messages associated with the selected Group.

Note that if you have the **Event Logs Profile** enabled on the **Profiles** tab of the **Options** dialog, the **Event Log Layout** automatically displays in the **Analysis Grid** viewer after you load data into Message Analyzer from a *.evtx file, as described in the Working With Message Analyzer Profiles topic. However, if you want to display the **Event Viewer Layout** for the **Grouping** viewer, you will need to manually select the **Default** item in the **Grouping** drop-down list that is accessible from the **New Viewer** drop-down list. Otherwise, if the **Event Logs Profile** is not enabled, you will need to manually select the **Event Viewer Layout** in the **Grouping** drop-down list, rather than the **Default** item, to use this **Layout**. The same is true of the **Chart** viewer **Layout** that is also part of the **Event Logs Profile**.

For more information about analyzing event logs with Message Analyzer, where such analysis is enhanced by complementary **Layouts** for the **Analysis Grid** and **Chart** viewers, see the **Event Logs Profile** in the table of the previously indicated topic.

- **Perfmon Log (.blg)** — this grouping **Layout** contains a top-level **Machine** Group with successively nested **Instance** and **Counter** Groups. The **Machine** Group identifies the computer/s on which the Performance Monitor data was collected; the **Instance** Group identifies the counter instances detected in the log; and the **Counter** Group identifies the Performance Monitor counters that collected the data.

  Note that if you have the **Perfmon Logs Profile** enabled on the **Profiles** tab of the **Options** dialog, the **Perfmon Log Layout** automatically displays in the **Analysis Grid** viewer and the appropriate **Perfmon Log Layout** automatically displays in the **Grouping** viewer after you load data into Message Analyzer from a *.blg file, as described in the Working With Message Analyzer Profiles topic. However, if the **Perfmon Logs Profile** is not enabled, you will need to manually select the **Perfmon Log Layout** in the **Grouping** drop-down list to use this **Layout**. The same is true of the **Chart** viewer **Layout** that is also part of the **Perfmon Logs Profile**.

  See the **Perfmon Logs Profile** in the table of the previously indicated topic for other details.

- **Protocol/Module Summary** — this grouping **Layout** provides a **Module** Group and a nested **Type** Group that enables you to explore data at a high-level. This layout summarizes your data so that you can analyze traffic volumes per protocol or module across a set of

messages, while the **Type** Group organizes messages by the type of message that a protocol or module issued.

For example, when you display this **Layout**, you might have multiple **Type** nodes nested under a particular **Module** parent Group, where each **Type** node indicates a different kind of message along with the traffic volume associated with each one. This **Layout** configuration is particularly useful with a protocol such as SMB that has many different message types.

In addition, this grouping **Layout** is intended to work with the **Raw Text Log Layout** of the **Analysis Grid** viewer to create an interactive analysis environment.

- **HTTP** category
  - **Fiddler Grouping** — this grouping **Layout** provides a top-level **SessionFlags.x-ProcessInfo** Group with a nested **Uri.Host** Group underneath it. The **Layout** isolates data from a Fiddler trace into groups, where you can view the message volume that is associated with each top-level process name and ID group (**SessionFlags.x-ProcessInfo**), the hosts that handled each request as indicated in the nested **Uri.Host** Group under a particular process name and ID Group, along with the number of messages (**Messages** column) associated with each host Group. By selecting either of these Groups, you can isolate the corresponding messages in the **Analysis Grid** viewer for assessment of message details.

    This grouping **Layout** is intended to work with the **Fiddler SAZ Layout** of the **Analysis Grid** viewer to create an interactive analysis environment. You will be able to correlate the data most effectively if you have this **Analysis Grid** viewer **Layout** displayed. For example, this **Layout** enables you to perform deep HTTP analysis of Fiddler traces with the data in the **StatusCode**, **Method**, **Uri.AbsPath**, **Uri, PayloadLength**, **Cache-Control**, **ContentType**, and **Payload** columns of the **Layout**. By selecting any particular Group, you can immediately expose and correlate all this data, including **Details**, for specific messages associated with the selected Group.

    Note that if you have the **Fiddler Traces Profile** enabled on the **Profiles** tab of the **Options** dialog, the **Fiddler SAZ Layout** automatically displays in the **Analysis Grid** viewer after you load data into Message Analyzer from a *.saz file, as described in the Working With Message Analyzer Profiles topic. However, if you want to display the **Fiddler Grouping Layout** for the **Grouping** viewer, you will need to manually select the **Default** item in the **Grouping** drop-down list that is accessible from the **New Viewer** drop-down list. Otherwise, if the **Fiddler Traces Profile** is not enabled, you will need to manually select the **Fiddler Grouping Layout** in the **Grouping** drop-down list, rather than the **Default** item, to use this **Layout**. The same is true of the **Chart** viewer **Layout** that is also part of the **Fiddler Traces Profile**.

    For more information analyzing Fiddler traces with Message Analyzer, where such analysis is enhanced by complementary **Layouts** for the **Analysis Grid** and **Chart** viewers, see the **Fiddler Traces Profile** in the table of the previously indicated topic.

  - **IIS** — this grouping **Layout** provides a top-level client IP address (**c_ip**) Group along with a nested server port (**s_port**) Group underneath it. The **Layout** enables you to isolate the client IP addresses that made requests to an IIS server, the IIS server ports that received the requests, and the query message volume sent to the server for each client IP address.

    This grouping **Layout** is intended to work with the **IIS Layout** of the **Analysis Grid** viewer to create an interactive analysis environment. You will be able to correlate the data most effectively if you have this **Analysis Grid** viewer **Layout** displayed. For example, this **Layout** enables you to analyze client requests for resources from an IIS site, along with other supporting data through the **s_sitename**, **cs_username**, **cs_method**, **cs_uri_stem**, **cs_uri_query**, and **csUser_agent** columns of the **Layout**. By selecting Groups in the

Grouping viewer, you can immediately expose and correlate all this data, including **Details**, for specific messages associated with a selected Group.

If you have the **IIS Logs Profile** enabled on the **Profiles** tab of the **Options** dialog, the **IIS Layout** automatically displays in the **Analysis Grid** viewer after you load data into Message Analyzer from an IIS.log file, as described in the Working With Message Analyzer Profiles topic. However, if you want to display the **IIS Layout** for the **Grouping** viewer, you will need to manually select the **Default** item in the **Grouping** drop-down list that is accessible from the **New Viewer** drop-down list. Otherwise, if the **IIS Logs Profile** is not enabled, you will need to manually select the **IIS Layout** in the **Grouping** drop-down list, rather than the **Default** item, to use this **Layout**. The same is true of the **Chart** viewer **Layout** that is also part of the **IIS Logs Profile**.

For more information analyzing IIS log data with Message Analyzer, where such analysis is enhanced by complementary **Layouts** for the **Analysis Grid** and **Chart** viewers, see the **IIS Logs Profile** in the table of the previously indicated topic.

- **File Sharing** category
  - **File Sharing SMB/SMB2** — this grouping **Layout** provides a top-level **SessionIdName** Group, along with a nested **TreeIdName** Group, which in turn contains a nested **FileName** Group. This Group configuration enables you to isolate messages and traffic volumes associated with SMB file sharing operations for different **TreeIds** that uniquely identify shares accessed during SMB sessions, which are in turn identified by **SessionIds**. The **Layout** enables you to view the message volume per session, as distinguished by the **SessionIdName** Groups, among potentially multiple sessions over a single SMB connection. You can also view specific share connections (TreeIds) via the nested **TreeIdName** Groups along with the nested **FileName** Groups under each parent **TreeIdName** Group. At each group level of the nested configuration, the **Grouping** viewer enables you to examine traffic volumes and to interactively drive display of messages associated with any selected group into the **Analysis Grid** viewer for further investigation and assessment of message details.

    For example, by clicking a **FileName** Group node in the **Group Values** column of this **Layout**, you can display all the messages that comprised the SMB operations that transpired during share/file access, such as Request, QueryInfo, Create, Read, Write, Close, and so on. This configuration conveniently extracts and organizes all this information into a condensed hierarchical format that enables quick analysis of file sharing issues, which could otherwise be an overwhelming task when you are dealing with very large message sets.

    This grouping **Layout** is intended to work with the **SMB Flat Layout** of the **Analysis Grid** viewer to create an interactive analysis environment. You will be able to correlate the data most effectively if you have this **Analysis Grid** viewer **Layout** displayed. After it is displayed, you can click the **Flat Message List** button on the Filtering toolbar to remove the encapsulation of SMB or SMB2 request and response messages as Operation nodes and return the response messages to their original chronological order in a set of trace results, similar to the Network Monitor view. This **Analysis Grid** viewer **Layout** enables you to analyze SMB and SMB2 file operations with the data in the **Source**, **Destination**, **SessionIdName**, **TreeIdNameReference**, **FileNameReference**, **Header.MessageId**, and **Summary** columns of the **Layout**. By selecting Groups in the **Grouping** viewer, you can immediately expose and correlate all this data, including **Details**, for the messages that are associated with a selected Group.

    If you have one or more of the **File Sharing SMB Profiles** enabled on the **Profiles** tab of the **Options** dialog, the **SMB Flat Layout** automatically displays in the **Analysis Grid** viewer after you load data into Message Analyzer from a corresponding *.cap, *.etl, *.pcapng, or *.pcap

file containing SMB or SMB2 messages, for which a **Profile** is enabled, as described in the Working With Message Analyzer Profiles topic. However, if you want to display the **File Sharing SMB/SMB2 Layout** for the **Grouping** viewer, you will need to manually select the **Default** item in the **Grouping** drop-down list that is accessible from the **New Viewer** drop-down list. Otherwise, if a **File Sharing SMB Profile** is not enabled, you will need to manually select the **File Sharing SMB/SMB2 Layout** in the **Grouping** drop-down list, rather than the **Default** item, to use this **Layout**. The same is true of the **Chart** viewer **Layout** that is also part of the **File Sharing SMB Profile**.

For more information about analyzing SMB and SMB2 messages with Message Analyzer, where such analysis is enhanced by complementary **Layouts** for the **Analysis Grid** and **Chart** viewers, see the **File Sharing SMB Profile** in the table of the previously indicated topic.

- **SysLog** — this grouping **Layout** provides a top-level Samba **level** Group, along with a nested **function** Group, which in turn contains a nested **source_file** Group. This Group configuration enables you to organize messages and traffic volumes based on the Samba debug **level**, the Samba **function/s** that wrote the Samba log entry, and the Samba **source_file** that contains the **function**. This **Layout** enables you to prioritize your investigation based on the level values, which is a good starting point from where you can determine, in a hierarchical manner, the functions and source code that is associated with the most critical debug levels in a SambaSysLog.

  This grouping **Layout** is intended to work with the **SysLog Layout** of the **Analysis Grid** viewer to create an interactive analysis environment. You will be able to correlate the data most effectively if you have this **Analysis Grid** viewer **Layout** displayed. With this **Layout**, you can analyze SambaSysLog data with the data in the **level**, **source_file**, **file_line**, **function**, **content**, and **Summary** columns of the **Layout**. By selecting Groups in the **Grouping** viewer, you can immediately expose and correlate all this data, including **Details**, for the messages that are associated with a selected Group.

  If you have the **Samba Logs Profile** enabled on the **Profiles** tab of the **Options** dialog, the **SysLog Layout** automatically displays in the **Analysis Grid** viewer after you load data into Message Analyzer from a SambaSysLog.log file, as described in the Working With Message Analyzer Profiles topic. However, if you want to display the **SysLog Layout** for the **Grouping** viewer, you will need to manually select the **Default** item in the **Grouping** drop-down list that is accessible from the **New Viewer** drop-down list. Otherwise, if the **Samba Logs Profile** is not enabled, you will need to manually select the **SysLog Layout** in the **Grouping** drop-down list, rather than the **Default** item, to use this **Layout**. The same is true of the **Chart** viewer **Layout** that is also part of the **Samba Logs Profile**.

  For more information analyzing SambaSysLog data with Message Analyzer, where such analysis is enhanced by complementary **Layouts** for the **Analysis Grid** and **Chart** viewers, see the **Samba Logs Profile** in the table of the previously indicated topic.

- **Netlogon** category
  - **Netlogon Group by Message Type** — this grouping **Layout** provides a single **msgtype** Group that enables you to isolate log entries in a Netlogon.log file based on different types of Netlogon messages. It also specifies the number of messages associated with each type. For example, message types from a Netlogon log can include CHANGELOG, CRITICAL, DNS, DOMAIN, LOGON, MAILSLOT, PERF, and so on. This simple Group configuration enables you to view which Netlogon message types exist in a Netlogon log, along with the message volume associated with each type, so that you can focus your analysis on specific messages that could potentially expose critical issues.

    This grouping **Layout** is intended to work with the **Netlogon Layout** of the **Analysis Grid**

viewer to create an interactive analysis environment. You will be able to correlate the data most effectively if you have this **Analysis Grid** viewer **Layout** displayed. For example, by selecting **Group Values** for a particular **msgtype** in the **Netlogon Group by Message Type Layout** of the **Grouping** viewer, you can interactively drive the display of only the messages of a selected Netlogon message type into the **Analysis Grid** viewer for further analysis. This means that you could select the CRITICAL message type in the **Group Values** column of the **Grouping** viewer and then only the messages of this type in your Netlogon log will display in the **Analysis Grid** viewer. With these messages isolated, you might then create a **Filter** in the Filtering toolbar text box to search for strings in the **Summary** column that indicate errors or failures, for example, a Filter such as the following: `*Summary contains "error"`.

If you have the **Netlogon Logs Profile** enabled on the **Profiles** tab of the **Options** dialog, the **Netlogon Layout** automatically displays in the **Analysis Grid** viewer after you load data into Message Analyzer from a Netlogon.log file, as described in the Working With Message Analyzer Profiles topic. However, if you want to display the **Netlogon Group by Message Type Layout** for the **Grouping** viewer, you will need to manually select the **Default** item in the **Grouping** drop-down list that is accessible from the **New Viewer** drop-down list. Otherwise, if the **Netlogon Logs Profile** is not enabled, you will need to manually select the **Netlogon Group by Message Type Layout** in the **Grouping** drop-down list, rather than the **Default** item, to use this **Layout**. The same is true of the **Chart** viewer **Layout** that is also part of the **Netlogon Logs Profile**.

For more information about analyzing Netlogon.log data with Message Analyzer, where such analysis is enhanced by complementary **Layouts** for the **Analysis Grid** and **Chart** viewers, see the **Netlogon Logs Profile** in the table of the previously indicated topic.

- **Network** category
  - **Network Address and Ports** — this **Layout** organizes trace results data into the different IPv4/IPv6 conversations at top-level in the **Network** groups, which display data in the **Group Values** column of the layout. Nested under this group are one or more **Transport** message Group nodes (see the TCP **Transport** field under the **TCP.Segment** node in **Field Chooser**). The **Transport** message nodes represent the different TCP/UDP ports on the source and destination computers that carried the IPv4/IPv6 conversations shown in the **Network** groups, where the messages associated with each **Transport** node are all in the same parent IPv4/IPv6 conversation.
  - **Process Name and Conversations** — this **Layout** isolates traffic into four Groups with the latter three successively nested under a parent **ProcessName** Group. In order, the nested groups consist of **ProcessName**, **ProcessId**, **Network**, and **Transport**, which together create a grouping configuration that isolates the TCP or UDP ports that carried IPv4/IPv6 conversations for various process IDs, each of which is associated with a specific process name, across a set of messages.

    Note that Message Analyzer can natively detect **ProcessName** and **ProcessId** information in a set of trace results and display it in this **Layout**. The input file types in which Message Analyzer can detect this information consist of *.etl, *.cap, and *.matp files. The retrieval of this information enables quick identification and correlation of process and conversation data. If you use this **Layout** with other file formats, Message Analyzer might be able to retrieve process names and process IDs for outgoing traffic, but might not be unable to obtain these for incoming messages. When this is the case, the **ProcessName** group values will duplicate the **ProcessId** Group values. The following describes how the Groups function in this **Layout**:

    - **ProcessName** — one or more **ProcessName** Groups display as top-level parent nodes in this **View Layout** to expose a process name for each **ProcessId** that is

associated with the conversations and ports that are specified in the nested **Network** and **Transport** child Groups.

- **ProcessId** — one or more **ProcessId** message Group nodes are nested under the **ProcessName** Groups and display data in the **Group Values** column based on the ETW **ProcessId** field (see the **Etw.EtwProviderMsg.EventRecord.Header.ProcessId** field in **Field Chooser**). This Group aligns the focus of this **Layout** configuration on outgoing traffic only, as it organizes messages by the **ProcessId** field that is found at the ETW layer. For example, in an HTTP operation, Request (outgoing) messages are evaluated for **ProcessId** values at the ETW layer, rather than evaluating the Response messages (incoming), mainly because the Request messages of an HTTP operation are the first to be encountered under top-level operation nodes, as displayed in the **Analysis Grid** viewer. As a result, this sets the precedent for the remaining evaluations to focus on Request (outgoing) messages.
- **Network** — one or more **Network** message Group nodes are nested under the **ProcessId** Groups and display data in the **Group Values** column based on the IPv4/IPv6 **Network** field (selecting the IPv4.Datagram.Network field in **Field Chooser** will suffice). The **Network** message nodes represent the different IPv4/IPv6 conversations that took place between source and destination computers, where the messages associated with each **Network** node all have the same **ProcessId** at the ETW layer.
- **Transport** — one or more **Transport** message Group nodes are nested under the **Network** Groups and display data in the **Group Values** column based on the TCP **Transport** field (see the **TCP.Segment.Transport** field in **Field Chooser**). The **Transport** message nodes represent the different TCP/UDP ports on the source and destination computers that carried the IPv4/IPv6 conversations, where the messages associated with each **Transport** node are all in the same parent IPv4/IPv6 conversation, have an identical **ProcessId**, and have the same **ProcessName**.

- **TCP Deep Packet Analysis** — this **Layout** organizes trace results data into the different IPv4/IPv6 conversations at top-level in the **Network** groups, which display data in the **Group Values** column of the layout. Nested under this group are one or more **Transport** message Group nodes (see the TCP **Transport** field (see the **TCP.Segment.Transport** field in **Field Chooser**). The **Transport** message nodes represent the different TCP/UDP ports on the source and destination computers that carried the IPv4/IPv6 conversations shown in the **Network** groups, where the messages associated with each **Transport** node are all in the same parent IPv4/IPv6 conversation. The **SourcePort** group exposes the port in a conversation on the source computer where a particular conversation was initiated.

Manipulating Group Displays

The two main ways that you can manipulate the grouped data display in the **Grouping** viewer to enhance your data analysis perspectives are as follows:

- **Expose different analytical contexts** — as described earlier, to achieve a different perspective on the data that displays in the **Grouping** viewer, you can pivot the grouped display by dragging and dropping any Group label into a different position in the hierarchic group arrangement that appears below the **Grouping** viewer toolbar, just as you can do with grouped messages in the **Analysis Grid** viewer, as described in Using the Analysis Grid Group Feature. This action changes the Group nesting configuration. When you alter the orientation of nested Groups, Message Analyzer refilters and repopulates the data in the current **Grouping Layout** according to the modified Group nesting configuration that you create. With this feature, you can change the relationship in which Group data is

displayed and thereby expose different analytical contexts and data correlations.

> **NOTE**
>
> You can also remove a Group in the **Grouping** viewer, by clicking on the "**x**" mark in its Group label. Also note that you can save any modifications that you make as a custom **Layout** of your own design, as described in Editing a Built-In Layout.

- **Summarize Group data** — to obtain a data summary for any Group, you can collapse a *particular* Group by right-clicking its label and selecting the **Collapse All Groups** context menu command. By collapsing a specific Group, you actually collapse the nested Group(s) under the data nodes of that specific Group. The resulting display provides a concisely summarized view of the data nodes associated with the particular Group for which you selected the **Collapse All Groups** context command. Thereafter, you can re-expand the Group with the **Expand All Groups** command to expose all data in the original nested Group configuration again. You can then move on to another Group in the **Grouping** viewer toolbar and do the same thing to achieve similar results. This enables you to quickly summarize the data associated with each Group label for clarity of analysis. You also have the option to collapse or expand all Groups simultaneously, as described earlier, by clicking the **Collapse All** or **Expand All** buttons, respectively, on the **Grouping** viewer toolbar.

Grouping Viewer Modes of Operation

The **Grouping** viewer has two modes of operation that provide different interactions with the **Analysis Grid** viewer. In the default **Filtering** mode, the selection of a node in the **Grouping** viewer filters and displays the messages in the **Analysis Grid** viewer that are associated with the selected node only and removes all others. In the **Selection** mode, the selection of a node in the **Grouping** viewer automatically selects/highlights the node messages in the **Analysis Grid** viewer. These interactions are best observed when the **Grouping** viewer and **Analysis Grid** viewer are docked side-by-side, as they are by default. The modes of operation for the **Grouping** viewer are explained in the following list.

- **Filtering Mode** — to enable this mode, click the **Filtering Mode** button (with the funnel-shaped icon) on the **Grouping** viewer toolbar.

> **TIP**
>
> To show the **Filtering Mode** button and other buttons on the toolbar with full text labels and icons, right-click the toolbar and select the **Show Labels and Icons** command in the context menu that appears.

In the **Filtering Mode**, selecting a node in the **Group Values** column of the **Grouping** viewer causes the messages associated with that node to be filtered in the **Analysis Grid** viewer such that only those messages display in the **Analysis Grid**. To remove the applied filtering, click the **Reset** button on the **Grouping** viewer toolbar.

> **NOTE**
>
> In the **Filtering Mode**, the **Selection** window does not track the messages that display in the **Analysis Grid** viewer.

- **Selection Mode** — to enable this mode, click the **Selection Mode** button (with the grid-shaped icon) on the **Grouping** viewer toolbar. In this mode, selecting a node in the **Group Values** column of the **Grouping** viewer causes the message/s associated with that node to be selected and highlighted in the **Analysis Grid** viewer. You can also select multiple message nodes by holding down the `Ctrl` key on your keyboard as you make selections. You can undo selections one at a time in the same manner, or you

can undo all selections at once by clicking the **Reset** button on the **Grouping** viewer toolbar.

> **TIP**
> If you have the **Selection Tool Window** open as you are making different selections in the **Grouping** viewer, the **Selection** window keeps track of all your message selections. This tracking feature enables you to use the **Go back** arrow-button on the **Selection** window toolbar to scroll through previous selections you made in the **Grouping** viewer. When you reach the last message or group of messages in the selection collection, you can then use the **Go forward** arrow-button to incrementally advance through those selections in the opposite direction.
>
> This improvement to analysis capabilities allows you to backtrack to a previous message selection you may have forgotten or navigate to a previous selection if you lose focus to another viewer. The selection feature maintains context while providing quick and convenient access to messag es of interest that you selected for analysis purposes.

More Information
To learn more about how to use the **Selection** window, see the Selection Tool Window topic.

Locating Analysis Grid Messages in the Grouping Viewer

If you want to determine where any message in the **Analysis Grid** viewer is located within the Group configuration of the current **Grouping** viewer **Layout**, you can simply right-click a message of interest in the **Analysis Grid** viewer and select the **Find in Grouping Viewer** command from the context menu that appears. Message Analyzer then locates the Group in which the message exists and selects and highlights that Group. If the message happens to be located in a collapsed parent Group, Message Analyzer will expand that Group to show the selection. This feature is best observed when the **Analysis Grid** and **Grouping** viewers are docked side-by-side, as they are by default.

After you locate an **Analysis Grid** viewer message in a particular Group, you can analyze it in the context of other related messages in the Group. You can also use the keyboard combination `Ctrl+Click` on successive Group nodes to aggregate additional related messages in the **Analysis Grid** viewer. If you have the **Selection** mode set in the **Grouping** viewer when you do this, the aggregated messages are all selected and highlighted in the **Analysis Grid** viewer to provide an instant view of interrelated messages, all in Group context. If you have the **Filtering** mode set, then Message Analyzer filters the selected messages such that only the aggregated messages display in the **Analysis Grid** viewer, thus providing a concise view of related messages in Group context for enhanced analysis perspective. These capabilities can also address a typical scenario where you are trying to correlate messages from different data sources, for example, a log file and live trace results, so that you can locate and analyze interrelated messages from such sources.

Grouping Viewer Display Features

The **Grouping** viewer display has a toolbar with various buttons decorated with icons in the upper section of the viewer, a row of movable Group labels below the toolbar, and a data grid underneath the labels. The toolbar buttons provide the following commands for manipulating the message Groups that display in the data grid section of the viewer:

- **Collapse All** — enables you to collapse all Group nodes to show the top-level Group only, as identified by the Group label to the far left in the labels row.
- **Expand All** — enables you to expand all Group nodes to expose the message data in the top-level and all nested Groups.
- **Add Groupings** — displays the **Field Chooser Tool Window**, if it is not already displayed; otherwise, it sets the **Field Chooser** as the in-focus window. Enables you to locate message fields that you want to add as a Group to the current Grouping **Layout**. To add a new Group, right-click

such a field in the **Field Chooser** and then select the **Add as Grouping** command in the context menu that appears.

- **Layout** — a drop-down list that enables you to select built-in **Layouts** from the **Message Analyzer Grouping View Layouts** asset collection Library, or any custom **Layouts** of your own, to create different grouped displays that augment your analysis capabilities.
- **Selection Mode** — causes node selection in the **Grouping** viewer to display/highlight messages in the **Analysis Grid** viewer that are associated with the selected node.
- **Filtering Mode** — causes node selection in the **Grouping** viewer to filter/isolate messages in the **Analysis Grid** viewer that are associated with the selected node.
- **Reset** — either removes the filtering or unselects all messages selected in the **Analysis Grid** viewer, depending on the current mode of the **Grouping** viewer.
- Toolbar context menu commands for the tools display format:
  - **Show Default Layout**
  - **Show Icons Only**
  - **Show Labels Only**
  - **Show Labels and Icons**

The data grid area of the **Grouping** viewer contains the following columns by default:

- **Group Values** — provides the data for each specified Group node in the nested grouping configuration.
- **Messages** — provides the number of messages associated with each Group node.

Context Menu Commands The **Grouping** viewer provides a context menu that displays the following commands when you right-click a row of data under a particular column:

- **Copy Selected Rows** — this command displays in the context menu irrespective of the column under which you right-click.
- **Copy 'Group Values'** — this command displays in the context menu when you right-click a Group under the **Group Values** column. Enables you to copy the Group name to the Clipboard.
- **Copy 'Messages'** — this command displays in the context menu when you right-click a Group under the **Messages** column. Enables you to copy the number of messages under a particular Group to the Clipboard.

Adding New Groups

You have the capability to add more groups to the view layout of the currently displayed **Grouping** viewer. To do this, click the orange-colored **Add Grouping** icon on the **Grouping** viewer toolbar to display the **Field Chooser Tool Window**. When you click this icon, the **Field Chooser** window either docks and displays in its default location if it was not already displayed, or receives the focus if it was already displayed. In either case, you must use the **Field Chooser** to locate a relevant message field to add to the current Grouping **Layout**. When you locate a message field that you want to add as a new Group, right-click it and select the **Add as Grouping** command from the context menu that appears. As a result, the new Group is nested at the *lowest* level in the **Grouping** viewer display, although you can alter the location of any group by dragging it to a different position in the hierarchic group arrangement. If you do this, Message Analyzer will refilter and reorganize the grouped data according to the new nesting configuration that you created. This can provide unique perspectives for data analysis.

TIP

When adding new groups to the **Grouping** viewer, you may want to think carefully about which fields you are adding, how they are related to the message data you are working with, what data you would like to extract and expose from your message set, and how your choice of fields will augment the analysis perspective you are attempting to achieve. To this end, you can assess some of the built-in **Layouts** for both the **Grouping** and **Analysis Grid** viewers to get a sense of why certain fields were chosen for the impact they make on streamlining the analysis process.

Note that you can also add Groups to the current Grouping **Layout** by right-clicking a field in the **Details Tool Window** and then selecting the **Add '<fieldName>' as Grouping** command from the context menu that appears. The placeholder *fieldName* in this command is the **Name** of the field that you right-click in **Details**. However, note that you must have the **Grouping** viewer in focus for the group to be added to it.

Editing a Built-In Layout

If you find that a particular **Grouping** viewer **Layout** works well for you but requires some tweaking for your environment, you can customize that **Layout** by adding other message field Groups to it, removing Groups, and/or reorganizing the Group nesting configuration. After modifying the **Layout**, you also have the option to save it under a different name so it will be available for selection going forward. Note that you can also save your custom **Layout** as the default by clicking the **Save Current as Default User Layout** command from the **Layout** drop-down list on the **Grouping** viewer toolbar. To customize a **Grouping** viewer **Layout**, you can use the **Field Chooser** to locate and add other message fields to the **Layout** to create additional groupings of data. To do so, you simply locate the field you want to add in the **Field Chooser**, right-click it, and then select the **Add as Grouping** command from the context menu that appears, as previously described. You can then select the **Save Current Layout As...** command in the **Layout** drop-down list to display the **Edit Item** dialog, from where you can specify a **Name**, **Description**, and **Category** for your new **Layout**. When complete, you can **Save** it to the **Message Analyzer Grouping View Layouts** asset collection Library. Thereafter, whenever you start Message Analyzer and run a Data Retrieval Session or a Live Trace Session, your custom **Layout** will be accessible from the **Layout** drop-down list on the **Grouping** viewer toolbar.

TIP

When you are editing a **Grouping** viewer **Layout**, you also have the option to add fields to the **Layout** from the **Details Tool Window**, as described earlier in Adding New Groups.

Creating a Grouping Viewer Layout Template

You will also use the **Field Chooser** to create your own **Grouping** viewer **Layouts**, in a manner that is similar to the way you use the **Group** command in the **Analysis Grid** viewer, as described in Using the Analysis Grid Group Feature. To create a new **Layout** for the **Grouping** viewer, you might start by creating and saving a blank template that you can later display whenever you want to configure message field Groups for a new **Grouping** viewer **Layout**. To do this, you can perform the following general workflow; the procedure assumes you already have the **Grouping** viewer displayed.

1. Remove all the existing groups from the currently displayed **Layout** in the **Grouping** viewer by clicking the **X** in each group label.
2. Click the **Layout** drop-down list on the **Grouping** viewer toolbar and select the **Save Current Layout As...** command to display the **Edit Item** dialog.
3. Specify a **Name** and **Category** for the layout, for example, "Blank Template" and "Templates",

respectively. You can also specify an optional **Description** to indicate that the layout is a template.

4. Click **Save** when complete.

Thereafter, when you want to create a new **Layout** for the **Grouping** viewer, perform the following steps.

1. From the **Grouping** viewer **Layout** drop-down list, select and display the blank **Layout** template you created in the previous procedure.
2. In **Field Chooser**, locate the field you want to add as a Group, right-click it, and then select the **Add as Grouping** command from the context menu that appears.

   Repeat step 2 for as many Groups that you want to create in your new **Layout**.

> **NOTE**
> The order in which you add message field groups determines the order in which Group(s) are nested. For example, the first message field you add to the **Grouping** viewer becomes the top-level Group node; the second message field that you add is nested under the top-level node as the second Group; the third field you add is nested under the second Group, and so on. Each Group that you add to the **Layout** is represented as a label above the data grid of the **Grouping** viewer and appears as soon as you add it.

3. When complete, save your new **Layout** configuration by selecting the **Save Current Layout As...** command in the **Grouping** viewer **Layout** drop-down list.

Thereafter, whenever you load your new **Layout**, it is automatically populated with data based on the message set of the in-focus **Analysis Grid** session viewer tab.

Managing Grouping View Layouts

The **Layout** drop-down list on the **Grouping** viewer toolbar provides commands that you can use to manage your **Layouts**, as follows.

- **Save Current Layout As...** — click this command to save any built-in **Layout** that you modified or any new **Layout** that you created. Selecting this command displays the **Edit Item** dialog, from where you can specify a **Name**, **Description**, and **Category** in which to place the **Layout**. Note that any custom **Layout** that you save will be placed in a top-level **My Items** category, under a subcategory that you specify.
- **Manage Layouts** — click this command to display a submenu with the following commands:
  - **Save Current As Default User Layout** — click this command to save the currently displaying **Grouping** viewer **Layout** as the default. Selecting this command will override the current default **Grouping** viewer **Layout** setting.
  - **Load Default User Layout** — click this command to display the **Layout** that you specified as the default with the **Save Current As Default User Layout** command.
  - **Restore Application Default Layout** — click this command to restore Message Analyzer's default Grouping viewer **Layout**, which contains a **Network** and **Transport** group.
  - **Manage...** — click this command to display the **Manage Grouping Layout** dialog, which has the common dialog format that Message Analyzer uses to manage all assets. From this dialog, you can use the **Import** function to retrieve Grouping **Layout** assets from a user-designated file share or other directory location for sharing purposes. Likewise, you can use the **Export** function to publish **Layout** assets to a share or other designated location to share your **Layouts** with others. You can also use the **Delete** command to remove any **Layouts** that exist

in the **My Items** category only. In addition, the **Manage Grouping Layouts** dialog has several context menu items that you can make use of after you right-click a **Layout**, as follows:

- **Edit** — click this command to display the **Edit Item** dialog, from where you can edit the **Name**, **Description**, and **Category** of a Grouping viewer **Layout**. This command is available only for **Layout** items under the top-level **My Items** category in the **Manage Grouping Layout** dialog.
- **Create a Copy** — click this command to create a copy of any existing **Layout** in any category. You then have the option to save it with a different **Name** and **Description**, and you can also save it in a different **Category**.
- **Delete** — click this command to delete the selected **Layout**. This command is available for **Layout** items under the top-level **My Items** category only, in the **Manage Grouping Layout** dialog.

See Also

ETW Framework Conceptual Tutorial
Working With Message Analyzer Profiles
Selection Tool Window
Using the Analysis Grid Group Feature

Pattern Match Viewer

2 minutes to read

Filtering is an important technique for isolating messages in a trace that meet specific filtering criteria. However, because the application of filtering is restricted to the domain of *individual* message values, it cannot expose the context or "sequence" in which events occur across the entire timeline of a trace. To enable sequences or patterns of events to be detected, Message Analyzer provides a pattern matching capability that can identify sequential message patterns in a *group* of messages. This pattern detection process is carried out by a Message Analyzer pattern matching engine that provides a fast and easy way for you to isolate and evaluate sequential patterns in your data. Pattern matching is a unique addition to the arsenal of Message Analyzer tools that you can use to analyze your message data.

Using Pattern Matching

You can use **Pattern** matching as a unique mode of analysis for trace results. In the hypothetical examples that follow, the pattern matching engine would identify the matched instances of an executing pattern definition within a trace, and when complete, report the messages that contain those occurrences, along with any additional information specified by the OPN behavior scenario design, such as captured field or property values. For example, you could use **Pattern** matching to accomplish any of the following:

- Identify virus patterns.
- Discover processes in a faulty state that return a consistent message pattern or sequence.
- Create TCP troubleshooting scenarios.
- Identify the request-response message pattern that is typical of HTTP, SMB, LDAP, DNS, and other protocols.
- Discover interesting areas to target for troubleshooting, rather than finding the specific cause of a problem.

Invoking Pattern Matching

In Message Analyzer, you can utilize pattern matching functionality by invoking the **Pattern Match** viewer from the locations described in Session Data Viewer Options. This includes specifying the **Pattern Match** viewer when starting a Data Retrieval Session, starting a Live Trace Session, or when opening it to assess trace results in an Analysis Session.

The figure that follows shows the results of executing the **TCP Three-Way Handshake** Pattern expression against a set of trace results. The **Pattern Match** viewer lists all the matches that were found in the **MATCHED INSTANCES** pane. In this pane, you will also find important statistics for TCP troubleshooting, such as TCP configuration settings that could expose potential TCP misconfiguration.

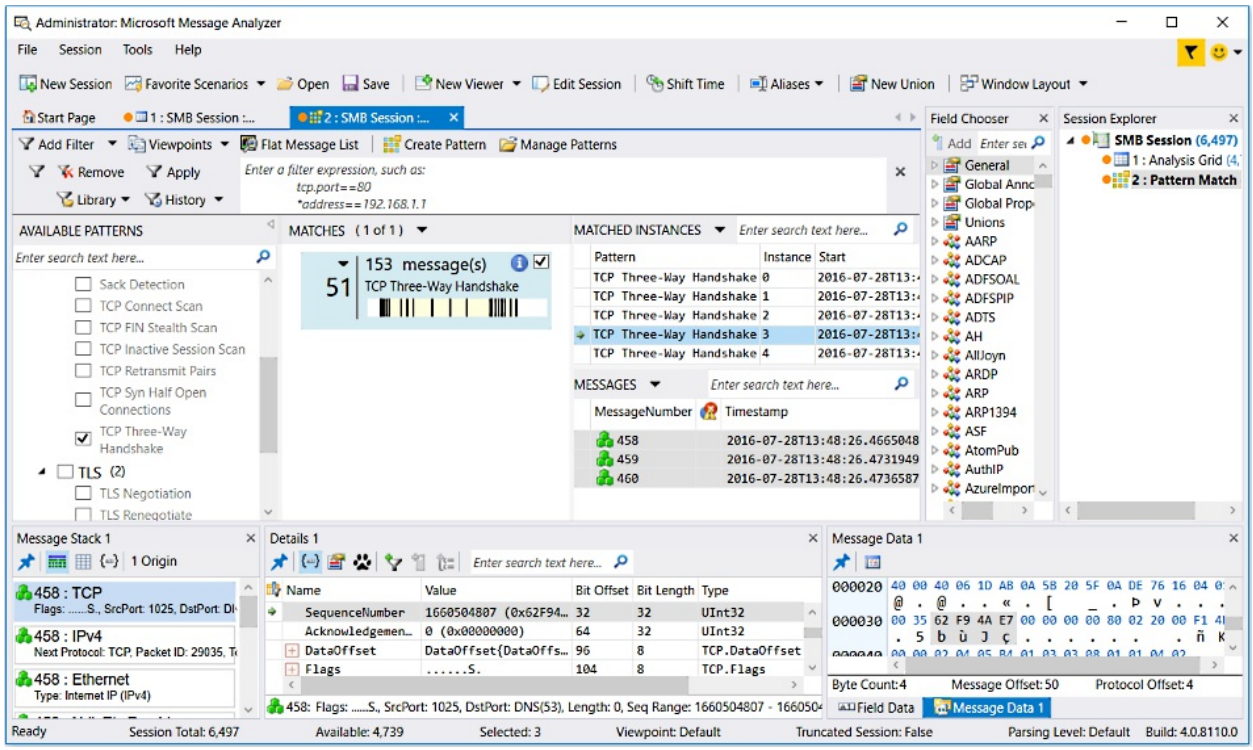

Figure 43: Pattern Match Viewer

What You Will Learn

In the following topics of this section, you will learn about matching message sequences by executing predefined **Pattern** expressions that are provided with the **Pattern Match** viewer. You will also learn about how to view matched instance data and how to create your own **Pattern** expressions:

Using the Pattern Match Viewer — learn how to use the **Pattern Match** viewer to execute predefined **Pattern** expressions and how to assess the resulting set of matched instances.

Understanding Message Pattern Matching — study an OPN code walkthrough for two of the predefined TCP **Pattern** expressions. These examples are intended to help you learn about **Pattern** expression construction and functionality, so that you can create your own.

Using the Pattern Editor — learn about how to build your own **Pattern** expressions in OPN, with or without the assistance of user interface (UI) automation.

Managing Pattern Expressions — learn how to manage **Pattern** expressions, which includes importing and exporting **Pattern** expressions for mutual sharing with others on your team.

Using the Pattern Match Viewer

11 minutes to read

This topic describes how to execute built-in **Pattern** expressions, how to view **Pattern** match data, and how to view messages that are associated with **Pattern** match data. To get started, you can launch the **Pattern Match** viewer against an existing set of trace results that are displayed in an **Analysis Grid** viewer instance. To do so, click the **New Viewer** drop-down list on the global Message Analyzer toolbar and select the **Pattern Match** item.

Executing Built-In Pattern Expressions

After you launch the **Pattern Match** viewer, you can execute a **Pattern** expression to view the results of pattern matching. You can do so by clicking one of the built-in **Pattern** expressions in the **AVAILABLE PATTERNS** pane of the **Pattern Match** viewer to execute that **Pattern** expression and return any pattern matches it finds.

> **NOTE**
>
> You can create one or more **Pattern** expressions of your own by clicking the **Create Pattern** button in the **Pattern Match** viewer toolbar. After you create a new **Pattern** expression, Message Analyzer adds it to the **My Items** category of the **AVAILABLE PATTERNS** list, from where you can select the expression to execute it against a set of trace results.

Currently, the built-in **Pattern** expressions that are available by default in every Message Analyzer installation are contained in the following category:

- **Network** category — the built-in **Pattern** expressions that are available in this category consist of the following:
  - **FTP Port Negotiate Failure** — locates negotiated FTP ports that are blocked with a Reset or fail to respond to a TCP SYN.
  - **FTP Port Negotiate Success** — searches for occurrences of FTP ports that are negotiated and then successfully set up as a TCP session.
  - **RPC Endpoint Mapper Failure** — searches for RPC ports that are negotiated through the RPC endpoint mapper, but blocked by the firewall or are ignored.
  - **RPC Endpoint Mapper Success** — locates occurrences of RPC port negotiation and of subsequent successful negotiation of a TCP session.
  - **Sack Detection** — searches for all selective acknowledgement (SACK) messages, which indirectly indicate the same network issues that cause TCP retransmits to occur, providing that SACK is enabled.
  - **TCP Connect Scan** — returns TCP sessions that are actively reset by the destination. Can be useful to find malware scans on the network.
  - **TCP FIN Stealth Scan** — searches for matches to TCP three-way handshakes that have no response, as an indication that a port is blocked or not listening. A significant number of matches could indicate network scanning is taking place.
  - **TCP Inactive Session Scan** — searches for TCP sessions where connection attempts have no response from an inactive port, which could be an indication of a security attack.
  - **TCP Retransmit Pairs** — enables you to identify pairs of retransmitted TCP messages with the same

sequence and acknowledgement numbers and an identical payload size, that occurred in the current set of messages.

- **TCP Syn Half Open Connections** — searches for TCP Syn/half open connections, such as a session that had a response where a port was opened but then Reset. Could indicate a port attack, for example, a denial of service (DoS) attack.
- **Three-Way Handshake** — enables you to isolate all three-way handshakes that occurred when setting up TCP connections, for both IPv4 and IPv6 transports, in the current set of messages. Also displays the approximate round trip time as the time delta between Syn messages and Syn Acknowledgement messages.

> **NOTE**
> In the latest Message Analyzer build, all of the TCP Pattern expressions are enabled to work with the **Microsoft-PEF-WFP-MessageProvider**, which has no IP/Network Layer.

- **TLS** category — the built-in **Pattern** expressions that are available in this category are described below.
  - **TLS Negotiation** — detects TLS sessions in which connection requests were initiated by a TLS client where the Session ID = 0, in other words, a cached Session ID was not used.

    The server then issues a Session ID to the client, which the client may or may not cache for use in subsequent connection requests. By caching the Session ID, the client can reuse the existing ID as a shortcut to facilitate subsequent connection requests, as required.
  - **TLS Renegotiate** — detects TLS sessions in which connection requests were initiated by a TLS client and the Session ID > 0, in other words, a cached Session ID was used.

    When a connection request for a TLS session is initiated by a TLS client and it reuses an existing Session ID, the client request is renegotiated to the TLS server with the use of a cached Session ID. However, during renegotiation, not all the TLS session information is represented on the wire. When this occurs, Message Analyzer's decryption expert is unable to decrypt these renegotiated sessions.

Therefore, the primary scenario in which you might execute these TLS **Pattern** expressions is when you run decryption against a set of trace results and you discover that decryption did not successfully occur, or a connection in which you were interested was not decrypted. By executing these **Pattern** expressions, you might learn that one or more connections in question were renegotiated. For example, you could determine this by examining which session connection requests were negotiated with an existing/cached Session ID and which ones were not.

Viewing Pattern Match Data

When you execute a **Pattern** expression, the results first display in the **MATCHES** pane of the **Pattern Match** viewer in a Matched pattern selector. To display the matched instances in the **MATCHED INSTANCES** pane of the **Pattern Match** viewer for the results of an executed **Pattern** expression, you can simply click the corresponding Matched pattern selector check box. This selector also has labels and controls that provide the following information or functions, respectively:

- **Matched instance count** — a label that indicates the number of matched instances that were returned by an executed **Pattern** expression.
- **Message count** — a label that indicates the total number of messages associated with all matched instances that were returned by an executed **Pattern** expression.
- **Information icon** — can link to a site that contains support information for the type of **Pattern** expression

that Message Analyzer executed. For example, the Information icon in the Matched pattern selector for the built-in **TCP Three-Way Handshake** pattern takes you to a site that provides more information about three-way handshake patterns. Note that you can specify a custom site to link to with the use of the **Remediation** feature when you are configuring a **Pattern** expression of your own.

- **Pattern expression name** — a label that indicates the name of the executed **Pattern** expression. Hover over this label with your mouse to display a tooltip with summary information for a particular results set, which can include a **Description** of the **Pattern** expression.
- **Check to include matches in the Instance list** check box — by placing a check mark in this check box or by removing it, you can alternately show or hide the results of an executed **Pattern** expression in the **MATCHED INSTANCES** list, respectively. This is useful when you are displaying multiple Matched pattern selectors that result from executing multiple **Pattern** expressions that returned data. By selecting and deselecting different Matched pattern selectors in the **MATCHES** pane, you can alternately display or hide the matched instances of any results set.

> **NOTE**
>
> If you are already displaying matched instances in the **MATCHED INSTANCES** pane of the **Pattern Match** viewer that result from executing a particular **Pattern** expression, you should unselect the corresponding Matched pattern selector for such **Pattern** expression results before you attempt to display the results of another **Pattern** expression execution by clicking a different Matched pattern selector. Otherwise, it could be more difficult to locate the latter data in the **MATCHED INSTANCES** pane, given that it is chronologically ordered and could be buried amidst other data.

- **Additional Options** drop-down menu — by clicking the black drop-down arrow on any Matched pattern selector, a menu appears with the following commands:
  - **Open Messages in Analysis Grid** — enables you to display the messages associated with a particular Matched pattern selector in the **Analysis Grid** viewer, for example, to examine message details or the message stack.
  - **Open Pattern Messages in Gantt** — enables you to display the messages associated with a particular Matched pattern selector in the **Gantt** viewer, for example, to examine the time window in which the matched instance messages occurred, along with source and destination IP address information.

Removing Pattern Match Data

To remove the results of an executed **Pattern** expression, simply click the check box for the executed **Pattern** expression in the **AVAILABLE PATTERNS** pane of the **Pattern Match** viewer to remove its check mark. This results in simultaneously removing the Matched pattern selector and any match instances data that are displaying in the **MATCHED INSTANCES** pane of the **Pattern Match** viewer for the particular **Pattern** expression that you deselected. You also have the option to simply unselect any Matched pattern selector to hide its associated matched instances data, until you select it again.

Viewing Matched Instance Message Data

The **MATCHED INSTANCES** pane displays all the instances that matched the **Pattern** expression that you executed, providing that you enable the corresponding Matched pattern selector. The **MATCHED INSTANCES** pane consists of the following elements:

- A down arrow that opens a drop-down list that contains the following commands:
  - **Create Bookmark** — this command enables you to create a bookmark for one or more selected matched instances in the **MATCHED INSTANCES** pane. All messages associated with each selected

matched instance will receive a bookmark.

For example, if you have the **Bookmarks Tool Window** displayed when you select two matched instances in the **MATCHED INSTANCES** pane and you execute the **Create Bookmark** command, all the messages that are associated with the selection will receive a bookmark and will appear in the **Bookmarks Tool Window** as a **Pattern Group**. Thereafter, if you click the book icon for the **Pattern Group** in the **Bookmarks Tool Window**, you will see all the messages that received a bookmark. If you then place the **Analysis Grid** viewer in focus (or dock it next to the **Pattern Match** viewer), you will be able to highlight each bookmarked **Pattern Group** message in the **Analysis Grid** viewer as you select them in the **Pattern Group** message list of the **Bookmarks Tool Window**.

- **Bookmark All Matches** — this command enables you to configure a bookmark for all matched instances currently displayed in the **MATCHED INSTANCE** pane. This results in creating bookmarks for all messages in each matched instance. If you have more than one Matched pattern selector enabled, the pattern groups in the **Bookmarks Tool Window** will have the same name as the executed **Pattern** expression, rather than being named a **Pattern Group**.
- **Open Messages in Analysis Grid** — this command enables you to display all messages in one or more selected matched instances from the **MATCHED INSTANCE** pane in a new instance of the **Analysis Grid** viewer for further assessment.
- **Open Pattern Messages in Gantt** — this command enables you to display all messages in any one selected matched instance from the **MATCHED INSTANCE** pane in a new instance of the **Gantt** viewer for further assessment.
- **Open Pattern Instances in Gantt** — this command enables you to display all messages in one or more selected matched instances from the **MATCHED INSTANCE** pane in a new instance of the **Gantt** viewer for further assessment.

- A search text box that enables you to locate a specified string in all rows of matched instance data where a match is found.
- Numerous columns of data that are relevant to the Pattern expression that you executed and for which you received results.

The **MESSAGES** pane exists below the **MATCHED INSTANCES** pane of the **Pattern Match** viewer. Each time you select a matched instance in the **MATCHED INSTANCES** pane of the **Pattern Match** viewer, the **MESSAGES** pane is populated with the messages that are associated with the matched instance that you selected. You can change the format in which these messages display by selecting any of the following format options in the **MESSAGES** drop-down list:

- **Show messages in a list** — the default selection that displays the messages associated with a matched instance in a simple list.
- **Show messages grouped by pattern** — select this option to display messages in the **MESSAGES** pane as an expandable node that is identified by an executed **Pattern** expression name. By clicking the expandable node, you can display all the messages associated with the executed **Pattern** expression. By default, the node is in the collapsed state.
- **Show messages grouped by sequential pattern match instance** — select this option to display messages in the **MESSAGES** pane as an expandable node that is identified by an executed **Pattern** expression name and the particular matched instance number that is selected in the **MATCHED INSTANCES** pane of the **Pattern Match** viewer. By clicking the expandable node, you can display all the messages associated with the executed **Pattern** expression. By default, the node is in the collapsed state.

Regardless of the message display format that you choose, whenever messages for a matched instance are

displaying in the **MESSAGES** pane of the **Pattern Match** viewer, you can right-click the matched instance messages and select one of the following context menu commands to display the messages:

- **Open All Messages** — opens all messages from a matched instance message group in a new instance of the **Analysis Grid** viewer, for example, to examine message details and stack information.
- **Open Selected Messages** — opens only the selected messages from a matched instance message group in a new instance of the **Analysis Grid** viewer, for further analysis.

Note that the **MESSAGES** pane also provides **MessageNumber**, **Timestamp**, and **Summary** data columns, and in some cases, other fields will be included as data columns for specific data that was detected by an executed **Pattern** expression. For example, the **TCP Three-Way Handshake Pattern** expression contains many additional data fields that can be useful for analyzing problems with the setup of TCP connections.

Also note that the drop-down list on the **MESSAGES** pane contains all the commands specified in the previous two lists.

See Also

Understanding Message Pattern Matching
Using the Pattern Editor

Understanding Message Pattern Matching

20 minutes to read

Message Analyzer enables you to process a set of trace results to retrieve groups of messages that meet the sequential pattern criteria of manually configured or built-in **Pattern** expressions. Whenever you execute a **Pattern** expression by selecting one in the **AVAILABLE PATTERNS** pane of the **Pattern Match** viewer, an API is accessed that calls into the pattern matching engine which starts the search for a pattern match based on the definition of the selected **Pattern** expression.

> **NOTE**
>
> The underlying technology that you use to create **Pattern** expressions is part of the full Open Protocol Notation (OPN) language, as described in the OPN Programming Guide.

Accessing Built-In Pattern Expressions

Message Analyzer provides you with the option to use several built-in OPN **Pattern** expressions that are included in every Message Analyzer installation by default. You can access the built-in expressions only after you open the **Pattern Match** viewer, where they are included as a list in the **AVAILABLE PATTERNS** pane of this viewer. You can run a **Pattern** expression against a set of trace results by simply placing a check mark in the check box to the left of the **Pattern** expression that you want to execute, for example, the **TCP Retransmit Pairs** expression. Message Analyzer then begins searching for patterns that match the **Pattern** expression criteria and initially displays any matches that are found in a Matched pattern selector that appears in the **MATCHES** pane of the **Pattern Match** viewer.

> **NOTE**
>
> Message Analyzer also provides an **Example Sequence Expression** for HTTP methods under the **My Items** category in the **AVAILABLE PATTERNS** pane of the **Pattern Match** viewer that you can edit however you want, as a practice development scenario.

> **IMPORTANT**
>
> Microsoft will continue to develop pattern matching capabilities and make them available either from the Message Analyzer Sharing Infrastructure, and/or as part of the installation package in ongoing Message Analyzer release versions. This will include additional built-in **Pattern** expressions and extended functionality for the pattern matching engine.

Understanding Pattern Expressions

This section provides a code walkthrough of two of the built-in **Pattern** expressions that are related to TCP operations. By understanding the concepts of these working **Pattern** expressions, you can apply this knowledge when creating expressions of your own. The built-in **Pattern** expressions that are presented as examples consist of the following:

- **TCP Retransmit Pairs** — the OPN behavior scenario that defines this **Pattern** expression is encapsulated in the following OPN code:

```
using TCP;
using IPv4;
using IPv6;
using Utility;

// Identifies pairs of TCP messages with same Sequence and Ack numbers, // the same payload size, and
the same
// IP source and destination addresses. Payload values must not be zero and KeepAlive messages
// are ignored. . .

virtual operation TCPRetrans
{
  uint SeqNum = seqNum != 0 ? seqNum : seqNum6;
  uint OrgMessageNumber = mn != 0 ? mn as uint : mn6 as  uint;
  uint DestMessageNumber = mnretrans != 0 ? mnretrans as uint : mnretrans6 as uint;
  (IPv4Address | IPv6Address) SourceAddress = sa != null ? sa : sa6;
  (IPv4Address | IPv6Address) DestinationAddress = da != null ? da : da6;

  override string ToString()
  {
    return "Frame: " + DestMessageNumber.ToString() + " Sequence Number " + SeqNum.ToString() +
    " is a duplicate of frame " + OrgMessageNumber.ToString();
  }
}
= backtrack(Segment{Payload.Count >= 1})
(
Segment{SequenceNumber is var seqNum, AcknowledgementNumber is var ack, !KeepAlive(value), Payload is
var pyl,
GetMessageNumber(value) is var mn}  \\IPv4.Datagram { SourceAddress is var sa, DestinationAddress is
var da}
  ->
  Segment{SequenceNumber == seqNum, AcknowledgementNumber == ack, Payload.Count == pyl.Count,
  GetMessageNumber(value) is var mnretrans}  \\IPv4.Datagram { SourceAddress.Octets == sa.Octets,
  DestinationAddress.Octets == da.Octets }
)
|
(
  Segment{SequenceNumber is var seqNum6, AcknowledgementNumber is var ack6, !KeepAlive(value), Payload
is var pyl6,
  GetMessageNumber(value) is var mn6}  \\IPv6.Datagram { SourceAddress is var sa6, DestinationAddress
is var da6}
  ->
  Segment{SequenceNumber == seqNum6, AcknowledgementNumber == ack6, Payload.Count == pyl6.Count,
  GetMessageNumber(value) is var mnretrans6}  \\IPv6.Datagram { SourceAddress.Octets == sa6.Octets,
  DestinationAddress.Octets == da6.Octets }
);

// The KeepAlive method projects the KeepAlive annotation into a
// boolean. The method returns only the ‘false’ case to ensure that
// KeepAlive messages are not evaluated. .
bool KeepAlive(Segment s)
{
  return (s#IsKeepAlive != nothing)? (s#IsKeepAlive as bool) : false;
}

uint GetMessageNumber(any message x)
{
  var origins = x.Origins;
  if (origins.Count == 0)
    return x#MessageNumber as uint;
  else
    return GetMessageNumber(origins[0]);
}
```

This OPN scenario finds all TCP messages that were retransmitted as lost TCP segments and presents them in the Message Analyzer **Pattern Match** viewer as matched instances that each contain a pair of

messages with an identical TCP SequenceNumber and AcknowledgementNumber. The results of this **Pattern** expression can provide an indication of network issues, for example, TCP retransmits are required because packets are being dropped by the network.

To retrieve these messages, this **Pattern** expression evaluates the input stream for the SequenceNumbers, AcknowledgementNumbers, and Payload values of TCP messages, along with the source and destination IP addresses to ensure a match. To simplify this walkthrough, only the IPv4 case is discussed, although the IPv6 case that follows the "|" operator (OR) in the code is functionally similar:

1. ***Virtual operation*** — following the using statements is a "virtual operation" that provides for definition of variables and exposure of properties that are used in **Pattern** expression evaluations; for example, to hold source and destination IPv4 or IPv6 address values to check for matches; and to hold other TCP field values that are used to create a custom output message for display in the Message Analyzer UI whenever this **Pattern** expression finds a match.

   Custom output message — as part of the virtual operation, the overridden ToString() method is used to create a custom output message based on values returned in the expression; it can also be any custom string value that you specify. The string value displays under the **Summary** column in the **MATCHED INSTANCES** pane of the **Pattern Match** viewer for each matched instance after the **Pattern** expression completes its evaluation.

2. ***Backtracking*** — the line of code containing the "backtrack" statement indicates that the evaluation should start with, and always backtrack to, the next TCP Segment that has a Payload count that is greater than or equal to one.

3. ***Segment evaluation*** — the line of code containing the first "Segment" statement retrieves the following entities from the first TCP message Segment in the trace results that has a Payload greater than or equal to one, and applies the indicated processing:

   - SequenceNumber — the variable "seqNum" is set to the numerical value of the first SequenceNumber.
   - AcknowledgementNumber — the variable "ack" is set to the numerical value of the first AcknowledgementNumber.
   - Payload — the variable "pyl" is set to the numerical value of the first Payload so that payload counts can be checked.
   - GetMessageNumber — a method that sets the variable "mn" to the current message number and returns it.

> **NOTE**
> Observe that the KeepAlive method explicitly rules out TCP KeepAlive messages for evaluation, because they are not considered to be TCP retransmits.

4. ***IP addresses*** — the line of code containing the "IPv4.Datagram" statement gets values that set the "sa" and "da" variables based on the values of the IPv4 SourceAddress and DestinationAddress, respectively. Note that the expression uses the double backslash characters "\\" to get at the IPv4 level of the TCP origins tree, as described in Browsing Message Origins.

5. ***Segment evaluation*** — the line of code containing the second "Segment" statement causes the expression to continue to evaluate the input stream until a TCP message Segment is found with:

   - A SequenceNumber that is identical to the numerical value currently set in the variable "seqNum".

- An AcknowledgementNumber that is identical to the numerical value currently set in “ack”.
- A Payload count value that is identical to the value currently set in “pyl.Count”.
- Source and destination IP addresses that match the last segment evaluation.

> **TIP**
> The line of code containing the “->” operator tells the expression to go forward evaluating messages until a match is found and to pass over all nonmatching messages.

Match found — if the SequenceNumber and AcknowledgementNumber are identical to those in the last segment evaluation, the payload count value is equal to “pyl.Count”, and the source and destination address are the same as the last segment evaluation, then the message is a TCP retransmit; in which case, the variable “mnretrans” is set to the retransmit message number which is returned by the GetMessageNumber method and passed to the virtual operation for output.

6. ***Output*** — when two TCP messages are detected that meet the indicated criteria of this **Pattern** expression, the appropriate variable values are then abstracted to the virtual operation, which extracts the values needed for the **Pattern** expression output to appear as a **Retransmit Pair** in the Message Analyzer **Pattern Match** viewer.
7. ***Continuing evaluations*** — the evaluation process backtracks to the next TCP message Segment that meets the indicated Payload requirements following the initial evaluation point, and the process is repeated until all messages in the input stream (the trace results) have been evaluated.

- **Three-Way Handshake** — the OPN behavior scenario that defines this **Pattern** expression is encapsulated in the following OPN code:

```
using TCP;
using IPv4;
using IPv6;
using Utility;

// Looks for a TCP 3-way handshake with an IPv4 or IPv6 transport.
// Returns the approximate response time that equals the time delta
// between the Syn -> SynAck.

// A Virtual Operation creates special messages where properties
// are exposed and display as columns, and also provide a friendly
// description. .

virtual operation TCP3Way
{
  TimeSpan AproxRTT = approxrtt != null ? approxrtt : approxrtt6;
  (IPv4Address | IPv6Address) Source = sa != null ? sa : sa6;
  (IPv4Address | IPv6Address) Destination = da != null ? da : da6;
  ushort SourcePort = sp != 0 ? sp : sp6;
  ushort DestinationPort = dp != 0 ? dp : dp6;
  int ServerMaxSegmentSize = SrvMaxSegSize !=0 ? SrvMaxSegSize : 0;
  int ServerScaleFactor = SrvScale !=0 ? SrvScale : 0;
  int ServerSackPermitted = SrvSackPerm !=0 ? SrvSackPerm : 0;
  int ClientMaxSegmentSize = CliMaxSegSize !=0 ? CliMaxSegSize : 0;
  int ClientScaleFactor = CliScale !=0 ? CliScale : 0;
  int ClientSackPermitted = CliSackPerm !=0 CliSackPerm : 0;

   override string ToString()
   {
    return "TCP Session";
   };
}
// scenario ThreeWayHandshake
```

```
// Note: former use of “scenario” is replaced by a virtual operation.

// Backtrack causes the sequence to start the evaluation with a TCP
// message that has it’s SYN flag set to True, and to always go back
// to evaluate the next message that has it’s SYN flag set to True,
// following the original evaluation. This causes all messages in
// the input stream (trace results) to be processed. .
= backtrack(Segment{Flags is TCP.Flags{SYN == true}})

// Check on IPv4 transport. .
(
// Look for TCP over IPv4 that is not adjacent in the stack; this
// allows for finding tunnel TCP sessions. The expression saves and
// remembers the Address/Port values and timestamp which is used
// for comparison in other parts of this expression. .

  \TCP.Segment{Flags is TCP.Flags{SYN == true}, SequenceNumber is var IPv4SeqNumRequest, SourcePort is
var sp,
  DestinationPort is var dp, GetTimestamp(value) is var starttime, value.TransportKey is var
transportV4,
  GetMaxSegSize(value) is var SrvMaxSegSize, GetSackPermitted(value) is var SrvSackPerm,
GetScaleFactor(value) is var SrvScale}
  \\IPv4.Datagram{SourceAddress is var sa, DestinationAddress is var da, value.NetworkKey is var
networkV4}
  ->
  \TCP.Segment{Flags is TCP.Flags{SYN == true, ACK==true}, SequenceNumber is var IPv4SeqNumResponse,
  AcknowledgementNumber == IPv4SeqNumRequest + 1, SourcePort == dp, DestinationPort == sp,
  GetTimestamp(value) - starttime is var approxrtt, value.TransportKey == transportV4,
GetMaxSegSize(value) = var CliMaxSegSize,
  GetSackPermitted(value) = var CliSackPerm, GetScaleFactor(value) = var CliScale}
  \\IPv4.Datagram{SourceAddress.Octets == da.Octets, DestinationAddress.Octets == sa.Octets,
value.NetworkKey == networkV4}
  ->
  \TCP.Segment{Flags is TCP.Flags{ACK==true}, SequenceNumber == IPv4SeqNumRequest + 1,
   AcknowledgementNumber == IPv4SeqNumResponse + 1, SourcePort == sp, DestinationPort == dp,
value.TransportKey == transportV4}
  \\IPv4.Datagram{SourceAddress.Octets == sa.Octets, DestinationAddress.Octets == da.Octets,
value.NetworkKey == networkV4}
)
|
//IPv6 transport version
(
  \TCP.Segment{Flags is TCP.Flags{SYN == true}, SequenceNumber is var SeqNumRequest, SourcePort is var
sp6,
  DestinationPort is var dp6, GetTimestamp(value) is var starttime6, value.TransportKey is var
transportV6}
  \\IPv6.Datagram{SourceAddress is var sa6, DestinationAddress is var da6, value.NetworkHashCode is var
networkV6}
  ->
  Segment{Flags is TCP.Flags{SYN == true, ACK==true}, SequenceNumber is var SeqNumResponse,
AcknowledgementNumber == SeqNumRequest + 1,
  SourcePort == dp6, DestinationPort == sp6,  GetTimestamp(value) - starttime6 is var approxrtt6,
value.TransportKey == transportV6}
   \\IPv6.Datagram{SourceAddress.Octets == da6.Octets, DestinationAddress.Octets == sa6.Octets,
value.NetworkHashCode == networkV6}
  ->
  Segment{Flags is TCP.Flags{ACK==true}, SequenceNumber == SeqNumRequest + 1,  AcknowledgementNumber ==
SeqNumResponse + 1, SourcePort == sp6,
  DestinationPort == dp6, value.TransportKey == transportV6}
   \\IPv6.Datagram{SourceAddress.Octets == sa6.Octets, DestinationAddress.Octets ==
da6.Octets,value.NetworkHashCode == networkV6}
);

// Property used to retrieve the timestamp which is only at the
// choke point message, which is the bottom one. .
DateTime GetTimestamp(any message x)
{
   var origins = x.Origins;
   if (origins.Count == 0)
```

```
     return x#Timestamp as DateTime;
   else
     return GetTimestamp(origins[0]);
}
// Gets the value of the MaxSegmentSize from TCP options for initial and // backtracked message
comparisons. .
int GetMaxSegSize(TCP.Segment x)
{
    if(x.Options != nothing && x.Options is array<TcpOption>){
        array<TcpOption> opts = x.Options as array<TcpOption>;
        foreach(TcpOption o in opts)
        {
            if(o is MaxSegmentSize)
            {
                MaxSegmentSize mss = o as MaxSegmentSize;
                return mss.MaxSegmentSize;
            }
        }
    }
    return 0;
}

// Gets the value of WindowsScaleFactor from TCP options for initial
// and backtracked message comparisons. .
int GetScaleFactor(TCP.Segment x)
{
    if(x.Options != nothing && x.Options is array<TcpOption>){
        array<TcpOption> opts = x.Options as array<TcpOption>;
        foreach(TcpOption o in opts)
        {
            if(o is TCP.WindowsScaleFactor)
            {
                WindowsScaleFactor s = o as WindowsScaleFactor;
                return s.ShiftCount;
            }
        }
    }
    return 0;
}

// Checks the SackPermitted value in TCP options for initial
// and backtracked message comparisons. .
int GetSackPermitted(TCP.Segment x)
{
    if(x.Options != nothing && x.Options is array<TcpOption>){
        array<TcpOption> opts = x.Options as array<TcpOption>;
        foreach(TcpOption o in opts)
        {
            if(o is TCP.SackPermitted)
            {
                SackPermitted s = o as SackPermitted;
                return s.Length;
            }
        }
    }
    return 0;
}
```

This OPN scenario finds all TCP messages that form a three-way communication pattern that is typical of setting up a TCP connection, where IPv4 or IPv6 are used as transports. The output of this **Pattern** expression can provide an overview of TCP connections which may indicate configuration problems that are impacting performance.

For reference and convenience, a table is included ahead that shows the pattern of SYN and ACK field values and relative SequenceNumber and AcknowledgementNumber value representations to illustrate

the signature of a three-way handshake pattern that successfully opened a TCP connection. To retrieve the TCP messages, this **Pattern** expression evaluates the input stream for SYN and ACK flag settings, SequenceNumbers, AcknowledgementNumbers, SourcePort, DestinationPort, round trip time, and the Source and Destination IP addresses to ensure a match.

Note that this **Pattern** expression also returns the values of various TCP fields to further enhance the analysis process, for example, the client and server **MaxSegmentSize**, **WindowsScaleFactor**, and **SackPermitted** TCP option values, which expose how these options were negotiated during the initial stages of a connection request. The values of these options are held by corresponding client and server variables, for example, the ClientScaleFactor and ServerScaleFactor vars. Note that the TCP options configuration might be the cause of TCP performance issues.

> **NOTE**
> To simplify this walkthrough, only the IPv4 case is discussed, although the IPv6 case that follows the "|" operator (OR) in the code is functionally similar:

Table 12. TCP Handshake Connection Negotiation

COMPUTER NODE	MESSAGE SENT	SYN FLAG VALUE	ACK FLAG VALUE	SEQUENCENUMBER	ACKNOWLEDGEMENTNUMBER
Sending	Connection Request	True	False	x	0
Receiving*	Request Acknowledgement	True	True	y	x+1
Sending	Sync Acknowledgement	False	True	x+1	y+1

The **TCP Three-Way Handshake Pattern** expression identifies this pattern across a set of trace results and returns them first to the **MATCHES** pane of the **Pattern Match** viewer as a summary button, which in turn you must click to display the full results set in the **MATCHED INSTANCES** pane of the viewer. To accomplish this, the code uses the following processing elements and constructs:

1. ***Virtual operation*** — following the using statements is a "virtual operation" that provides for definition of variables and exposure of properties that are used in **Pattern** expression evaluations; for example, to hold source and destination IPv4 or IPv6 address values to check for matches; a round trip time value; and to hold other TCP field values that are used in the evaluations such as SourcePort and DestinationPort. Other declarations in the virtual operation are for variables that hold the values of TCP **MaxSegmentSize**, **WindowsScaleFactor**, and **SackPermitted** options on the sending and receiving nodes. These variables correlate TCP option values in the returned **Pattern** expression results for analysis purposes.

   Custom output message — as part of the virtual operation, the overridden ToString() method is used to create a custom output message, which can be based on values returned in the expression, or it can be any custom string value that you specify. The string value displays in the **Summary** column of the **Pattern Match** viewer when you click a **Matches** pane summary button after the executing **Pattern** expression completes its evaluation.

2. ***Backtrack*** — the line of code that contains the "backtrack" statement tells the **Pattern** expression to always start with, or backtrack to, a TCP segment that has its SYN flag set to true, which is indicative of a TCP message that is requesting to open a connection on a target node.

3. ***Segment evaluation*** — the line of code containing the first "TCP.Segment" statement looks at the first TCP message that has its SYN flag set to true, which is characteristic of a connection request from a sending node, that can be *accepted* by a receiving node. The **Pattern** expression then sets the values of the following variables for comparison with other messages:

   - "IPv4SeqNumRequest" — is set to the SequenceNumber of this initial TCP request message.
   - "sp" — is set to the SourcePort value of this message to identify the source port involved in the IP conversation where the handshake is being set up.
   - "dp" — is set to the DestinationPort value of this message to identify the destination port involved in the IP conversation where the handshake is being set up.
   - "starttime" — is set to the return value of the GetTimestamp() method. The time value in this variable is subtracted from the return value of the GetTimeStamp() method in the next TCP.Segment section to calculate the approximate round trip time that is set in the variable "approxrtt" in that section of code.
   - "transportV4" — is set to the value of the **TCP.Segment.TransportKey** field. Provides a hash constant value that is the difference between the source and destination TCP ports currently under evaluation. Provides a faster comparison of the current TCP port pair.

   > **NOTE**
   > For further details, right-click the **TransportKey** field in **Field Chooser Tool Window** and select **Go To Definition** in the context menu that appears to show the OPN definition of this field. You can also add this field to the **Analysis Grid** viewer column layout to see what a typical value might look like. Also, right-click a field value and select the **Include Hex for Numeric Values** command in the context menu that appears, to view the hexadecimal value that is used as the hash constant.

   - "SrvMaxSegSize" — is set to the value of the TCP MaxSegmentSize option; the method GetMaxSegSize() is called to return the value.
   - "SrvSackPerm" — is set to the value of the TCP SackPermitted option; the method GetSackPermitted() is called to return the value.
   - "SrvScale" — is set to the value of the TCP WindowsScaleFactor option; the method GetScaleFactor() is called to return the value.
   - "sa" — is set to the IPv4 SourceAddress of this message to identify the source node that is involved in the IP conversation where the handshake is being set up.
   - "da" — is set to the IPv4 DestinationAddress of this message to identify the destination node that is involved in the IP conversation where the handshake is being set up.
   - "networkV4" — is set to the value of the **IPv4.Datagram.NetworkKey** field. Provides a hash constant value that is the difference between the source and destination IPv4 addresses under evaluation. Provides a faster comparison of the current IPv4 address pair.

4. ***Segment evaluation*** — the line of code containing the second "TCP.Segment" section looks for the next TCP message that has both its SYN flag and ACK flag set to true, which is the signature of an acknowledgement reply *issued* by a receiving node to indicate a successful connection, that is, along with appropriate TCP SequenceNumber and AcknowledgementNumber values. This line of code also captures the SequenceNumber of this message in the "IPv4SeqNumResponse" variable and detects whether the AcknowledgementNumber of this message is equal to the SequenceNumber of the first message sent, plus 1. This would indicate that the message is a component part of the handshake, as long as the message is in the same established IP conversation and port communication, as they would be in an associated ACK

response. To verify that this is the case, the **Pattern** expression checks the values of the following variables:

- "dp" — is checked against the current message SourcePort to see if the values are identical. Note that the source and destination ports switch for the first ACK message.
- "sp" — is checked against the current message DestinationPort to see if the values are identical.
- "approxrtt" — is calculated by subtracting the "starttime" value of the last "TCP.Segment" from the current return value of the GetTimeStamp() method. This value is specified in the **ApproxRTT** column of the **Pattern Match** viewer when you expand the results node after the **Pattern** expression completes its evaluation.
- "transportV4" — is set to the value of the **TCP.Segment.TransportKey** field, as previously described.
- "CliMaxSegSize" — is set to the value of the TCP MaxSegmentSize option; the method GetMaxSegSize() is called to return the value.
- "CliSackPerm" — is set to the value of the TCP SackPermitted option; the method GetSackPermitted() is called to return the value.
- "CliScale" — is set to the value of the TCP WindowsScaleFactor option; the method GetScaleFactor() is called to return the value.
- "da.Octets" — is checked against the current message IPv4 SourceAddress.Octets value to see if the values are identical.
- "sa.Octets" — is checked against the current IPv4 Destination Address.Octets value.
- "networkV4" — is set to the value of the **IPv4.Datagram.NetworkKey** field, as previously described.

5. ***Segment evaluation*** — the line of code containing the third "TCP.Segment" section looks for the next TCP message that has its ACK flag set to true, its SequenceNumber equal to the value of the first message's SequenceNumber plus 1 ("SeqNumRequest +1"), and an AcknowledgementNumber equal to the value of the second message's SequenceNumber plus 1 ("SeqNumResponse + 1"), while also ensuring that the message is in the same IP conversation and is using the same ports. This pattern is characteristic of a successful connection that is in turn acknowledged by the sending node that made the initial connection request, which can be *accepted* by the receiving node.
6. ***Output*** — as these patterns are met, Message Analyzer returns the three-way handshake messages to the **Pattern Match** viewer for display. To learn more about how to display the results of executing a **Pattern** expression, see Using the Pattern Match Viewer.
7. ***Continuing evaluations*** — the **Pattern** expression then backtracks to the next TCP segment (where the SYN flag = true) following the initial evaluation point and the process is repeated until all messages in the input stream (trace results) have been evaluated.

Writing Pattern Expressions

Message Analyzer enables you to write your own **Pattern** expression in "free form", meaning that you must develop the expression without any UI automation support such as you have from the **Quick** tab of the **Pattern Editor**. To write your own **Pattern** expression in "free form", the **Pattern Editor** provides the features you will need. However, you will most likely need to learn more about OPN, although you can review the code descriptions provided here to help you obtain a rudimentary understanding of how the language works for building **Pattern** expressions.

TIP

When writing **Pattern** matching scenarios, it might be helpful to make the assumption that conversations are initiated by the client (sender) and that Message Analyzer is running on the server, the (*receiving) node (see the previous table).

When writing **Pattern** matching scenarios, you can take advantage of a friendly method for creating **Patterns**, by using the **Quick** tab of the **Pattern Editor**. The simplest way to do this is to right-click a message of interest in the **Analysis Grid** viewer and select the **Create Pattern** command in the context menu that appears. This action will open the **Pattern Editor** pre-populated with information that derives from the selected message.

More Information

To learn more about the **Pattern Editor**, see Using the Pattern Editor.

To learn more about executing sequential **Pattern** expressions, see the TechNet Blog article Sequence Match View: Identifying Interesting Network Patterns.

To learn more about OPN, see the OPN Programming Guide.

Using the Pattern Editor

18 minutes to read

You can generate a sequential pattern matching definition as an OPN behavior scenario or a virtual operation. Note that the term "scenario" or "virtual operation" is used in the OPN code for various **Pattern** expressions to identify to the OPN Compiler the OPN entity type that is being referenced. A virtual operation enables you to abstract property values from pattern matching definition results to provide additional processing, such as writing custom output messages. A "scenario" does not provide this particular capability, but in all other respects, scenarios and virtual operations are identical.

The purpose of writing an OPN behavior scenario is to find the specific sequential message patterns that the scenario defines. Pattern matching is built upon a pattern matching engine that is called by OPN code via an API. This makes it possible for you to create your own custom **Pattern** expressions as OPN behavior scenarios that are subject to OPN compilation, which thereafter makes them accessible to the Message Analyzer Runtime parsing engine. You can create your own **Pattern** expressions by configuring one with the **Pattern Editor** dialog. Similarly, you can edit any **Pattern** expression with the features of this dialog.

Accessing the Pattern Editor

To create or edit a **Pattern** expression, you must use the **Pattern Editor**. You will first need to display data from a Live Trace Session or a saved file in the **Analysis Grid** viewer before you can work with the **Pattern Editor**. After you display a set of trace results, you can open the **Pattern Editor** by executing the **Create Pattern** command in either of the following locations, but with slightly different results:

- **Pattern Match** viewer — click the **Create Pattern** button on the toolbar above the **AVAILABLE PATTERNS** pane of the **Pattern Match** viewer. When you open the **Pattern Editor** from this location, it displays in a separate tab with a blank configuration, so that you can insert message types and other parameters to create pattern definition components.
- **Analysis Grid** viewer — the **Create Pattern** command appears as a context menu item that displays when you right-click one or more selected messages in the **Analysis Grid** viewer. When you open the **Pattern Editor** from this location, it is prepopulated with an initial configuration of message **Id**, **Type**, and **Summary**, for *each* selected input message. The initial configuration provides a starting point for **Pattern** expression building, where you can specify parameters associated with the selected messages to create your pattern definition components.

> **TIP**
> This method of starting **Pattern** expression configuration is very useful because it helps you learn how to write patterns, given that the initial OPN code is automatically created for you.

After you open the **Pattern Editor** dialog from either of these locations, you can start to build pattern definition components for your **Pattern** expression, as described in the section that follows.

Tools for Building a Pattern Expression

The **Pattern Editor** dialog enables you to build your own **Pattern** expressions based on pattern definition components that you configure for specific message types. For example, the dialog enables you to configure the patterning criteria for one or more messages, by specifying various fields, properties, methods, and flags from the **Field Chooser Tool Window** — and/or values to which any of these entities can be set for a particular message

field, along with relational, logical, or numeric operators that qualify how such fields and properties are manipulated to create one or more criteria sets. Each criteria set will then function as a particular component of a pattern definition. After you launch the **Pattern Editor** as described in Accessing the Pattern Editor, you have the option to create a **Pattern** expression in two different ways, as follows:

- **Quick** configuration — a tabbed space that enables you to take advantage of UI automation to configure simple pattern definition components. You will find the *quick* configuration space on the **Quick** tab of the **Pattern Editor**. Note that this tab is prepopulated with data that is derived from one or more messages, but only if you launched the dialog from the previously described **Analysis Grid** viewer context menu command.

Caution

If you are working with the UI automation feature on the **Quick** tab of the **Pattern Editor** and you click **Edit** on the toolbar of the **Free Form** tab, you will be unable to return to the **Quick** configuration tab.

- **Free Form** configuration — enables you to write OPN code by hand to configure pattern definition components. You will find the *free form* configuration space on the **Free Form** tab of the **Pattern Editor**. This option is mainly for users who are familiar with OPN coding and can create more complex **Pattern** expressions.

Unpopulated Configuration Mode

If you launched the **Pattern Editor** by clicking the **Create Pattern** button in the **Pattern Match** viewer, as indicated earlier, it opens to the **Quick** tab in unpopulated configuration mode. Thereafter, assuming that you elect to perform **Quick** configuration rather than **Free Form**, you can use the controls of the **Pattern Editor** to create pattern definition components by performing the following operations:

- **Specify the expression name** — enter a descriptive name for the **Pattern** expression by typing it into the **Name** text box at the top of the editor.
- **Specify an expression description** — enter a conceptual description of the **Pattern** expression by typing it into the **Description** text box.
- **Set the expression category** — click the **Category** combo box drop-down list to include your **Pattern** expression in an existing category, or create a new one by typing it in the combo box.
- **Remediation** — specify a website address that provides support information related to the **Pattern** expression you are creating. The site will be opened when you click the Information icon in the Matched pattern selector that displays in the **MATCHES** pane of the **Pattern Match** viewer, after Message Analyzer finds one or more matched instances for an executed **Pattern** expression.
- **Insert a message** — under the **CRITERIA** subtab, click the **Insert message** button in the lower section of the **Pattern Editor** to open the **Field Chooser** window, from where you can select a message type for a particular protocol or module. Click **Insert message** for each message type that you want to add.
- **Set backtracking** — select the **Backtrack to find each match** check box if you want the **Pattern** expression to backtrack to the next evaluation point following the previous one in a set of trace results, to continue applying the filtering criteria of your pattern definition to the remaining messages.

More Information

To learn more about backtracking, see the predefined **TCP Three-Way Handshake** expression walkthrough in [Understanding Message Pattern Matching] (understanding-message-pattern-matching.md).

- **Insert a message criteria clause** — click the **Insert Criteria** link to display a set of clause configuration controls that enable you to perform the following basic operations:
  - Specify message fields, properties, and so on.
  - Specify an arithmetic operator, such as "+", "-", "/", or "*".

- Specify a relational operator such as "==" or ">=", the logical inequality (!=) operator, and so on. You can also specify the **contains** or **assign** functions.

  Note The **contains** function enables you to determine whether a value exists in a particular field and the **assign** function enables you to assign a field (that you specified in either the first or second input box of a particular criteria set) to a variable name that you create (in the third input box of a criteria set). You can also assign a variable to the results of an arithmetic operation upon two fields specified in the first and second input boxes of a particular criteria set.

The resulting configuration that you specify with these controls will construct an OPN clause that acts as a filter expression with left- and right-hand side components, which subsequently constitutes a pattern definition component that forms a part of your overall **Pattern** expression. As a simple example of a criteria set for an HTTP.Request message type, you could specify the **Method** field in the first criteria box, leave the second box empty, specify the "==" relational operator in the second drop-down list, and then type "GET" in the third criteria box.

- **Set a field name** — click the ellipsis (**...**) control to the right of the first criteria input box to display the **Field Chooser**, from where you can select message fields, properties, methods, flags, and so on.
- **Set an operator** — specify an arithmetic operator in the drop-down list that separates the first two criteria input boxes. For example, you might want to add or subtract the values of two fields, the second of which you can specify in the next criteria input box, as described in the next bullet point.
- **Set another field name** — click the ellipsis to the right of the second criteria input box to specify another field name, for example, one that is of the same or related type as the one you specified in the first criteria input box. For instance, you could subtract the values of two fields and check the results against some constant or other value that you specify in the third criteria input box. You can also specify another field name in the third criteria input box rather than a value, since Message Analyzer will automatically obtain its value when executing the **Pattern** expression. Thus far, the filter expression for this criteria set would look similar to the following:
  ```
  field1Name - field2Name
  ```
- **Set another operator** — specify a logical or relational operator such as greater than (>) in the second operator drop-down list to establish how the results of the previous *fieldName* operands will be evaluated.
- **Specify a value or field** — enter a value in the third criteria input box, against which the *fieldName* operand results are to be evaluated. The resulting filter expression for this criteria set would look similar to the following:
  ```
  field1Name - field2Name > someValue
  ```

> **TIP**
> Note that Message Analyzer enables you to use **Timestamp** as a field in a Pattern Expression. **Timestamp** is a **Global Annotation** that is accessible from the **Field Chooser** window.

- **Create a Summary for matched instance display** — in the text box under the **SUMMARY** subtab, specify summary text that will appear under the **Summary** column in the **MATCHED INSTANCES** pane of the **Pattern Match** viewer. This can include the use of variable values that are specified in the **Pattern** expression.

NOTE

You can include additional columns of data for variable values in the **MATCHED INSTANCES** pane of the **Pattern Match** viewer, providing that you used the **assign** function in the **Pattern Editor** to assign the value of a field, or an operation on two fields, to a variable name that you create. To add columns of data based on one or more variables, you can select them under the **Column** heading that appears below the **SUMMARY** subtab text box. You also have the option to override the column (variable) names by specifying a different column name in the **Name Override** column beneath the **Summary** text box.

TIP

To review an example of how to use the **assign** function, see Example of Building a Simple Pattern Expression.

- **Specify the message evaluation direction** — on the **ADVANCED** subtab of the **Quick** tab, optionally place a check mark in the **Evaluate sequential message patterns using the oldest to newest message** check box. This will begin the evaluation from the captured message that has the latest **Timestamp** value in a set of trace results; otherwise, the default (this check box unselected) is to start the evaluation from the message that contains the earliest **Timestamp** value in a set of trace results.
- **Specify a Viewpoint** — on the **ADVANCED** subtab of the **Quick** tab, click the **Viewpoint** drop-down list to select a **Viewpoint** that will enable your **Pattern** expression to evaluate messages from the perspective of a particular protocol or other module, while filtering out all messages above the **Viewpoint**. This can ensure that your **Pattern** expression processes all the message types on which your pattern is focused.

More Information

To learn more about **Viewpoints**, see [Applying and Managing Viewpoints] (applying-and-managing-viewpoints.md).

Populated Configuration Mode

If you launched the **Pattern Editor** by right-clicking two or more messages in the **Analysis Grid** viewer and then selecting the **Create Pattern** command from the context menu that appears, the editor opens to the **Quick** tab in the populated configuration mode with message types for each selected input message already inserted. Thereafter, the configuration options that you can use to create pattern definition components consist of all those specified immediately above. An example of creating a **Pattern** expression in this mode is described in Example of Building a Simple Pattern Expression.

Free Form Configuration Mode

If you choose to create a **Pattern** expression without any UI automation support, you can click the **Free Form** tab of the **Pattern Editor** and enter OPN code that creates your pattern definitions. However, you should have some familiarity with OPN before you try this option, since OPN is a unique language with many different constructs, operators, specifiers, and semantic and syntactic representations. Also, once you select the **Free Form** tab and click **Edit**, you are prompted that you will be unable to return to the **Quick** tab again with the UI automation configuration capability for the current editing session.

Example of Building a Simple Pattern Expression

The steps that follow provide an example of how to build a simple **Pattern** expression. This **Pattern** expression locates pairs of TCP messages, where the first message has a particular TCP **SequenceNumber** and the second reflects the **NextSequenceNumber** of the TCP message that carries the next payload segment, which is defined in the expression as equal to the **PayloadLength** of the second message plus the **SequenceNumber** of the first message:

1. Open the **Pattern Editor** by right-clicking two TCP messages that are selected in the **Analysis Grid** viewer and then selecting the **Create Pattern** command in the context menu that appears.

   The **Pattern Editor** opens to the **Quick** tab with some preliminary data for the selected messages displayed.

2. Click **Insert Criteria** three times for each message, where each message is labeled with an **Id** that specifies **A** or **B**.

   Thereafter, three rows of criteria input controls are displayed for each TCP message.

3. Click the ellipsis (**...**) to the right of the first criteria input box to display the **Field Chooser** window, then double-click the **SequenceNumber** field in the TCP message hierarchy.

   The **SequenceNumber** field name appears in the first criteria input box.

4. Click the second drop-down list (to the left of the third criteria input box) and select the **assign** item; then create a variable name in the third criteria input box to hold the value of the **SequenceNumber** field, for example, "varSeqNumber".

   This variable will be used in the first criteria row that you insert for the second TCP message, where the value of "varSeqNumber" is added to the **PayloadLength** value so that the **Pattern** expression will find a message with a **NextSequenceNumber** field that is equal in value to this sum.

5. In the second row of criteria controls, click the ellipsis (**...**) to the right of the first criteria input box to display the **Field Chooser**, then double-click the **SourcePort** field to display it in the first criteria input box.

6. Using the method indicated in step 4, create a variable such as "varSourcePort" and **assign** it to the **SourcePort** field to hold its value.

7. Repeat the previous two steps to create a variable such as "varDestPort" to hold the value of the TCP **DestinationPort** field. Use the **assign** function to associate the field with the variable.

   These port values will be checked against port values in subsequent TCP messages that are evaluated by the second set of message criteria, to ensure that the **Pattern** expression processes only messages that are in the same conversation.

> **NOTE**
>
> Whenever you create a variable as the right-hand expression of a criteria set and assign the value of some field to it in the **Pattern Editor**, that variable will appear in **Field Chooser** under the **Sequence Variables** node, which displays in the context of the **Pattern Editor** only.

8. In the first row of criteria controls for the second TCP message, click the ellipsis (**...**) to display the **Field Chooser** and then double-click the **PayloadLength** field in the TCP message hierarchy to display this field name in the first criteria input box.

9. In the drop-down list immediately to the right of the first criteria input box, select the **+** item to configure the *addition* arithmetic operator.

10. In the criteria input box to the right of the drop-down where you selected the arithmetic operator, type the name of the first variable you created, for example, "varSeqNum".

11. Click the ellipsis to the right of the third criteria input box and then double-click the **NextSequenceNumber** field in **Field Chooser** to add its name to the third input box.

12. In the second and third row of criteria controls for the second TCP message, click the ellipsis to display

Field Chooser and then double-click the **DestinationPort** and **SourcePort** fields, respectively, to add these field names to the first criteria input boxes of the second and third row of criteria controls, respectively, for the second message.

When complete, the **Pattern Editor** dialog should reflect a configuration that is similar to the following:

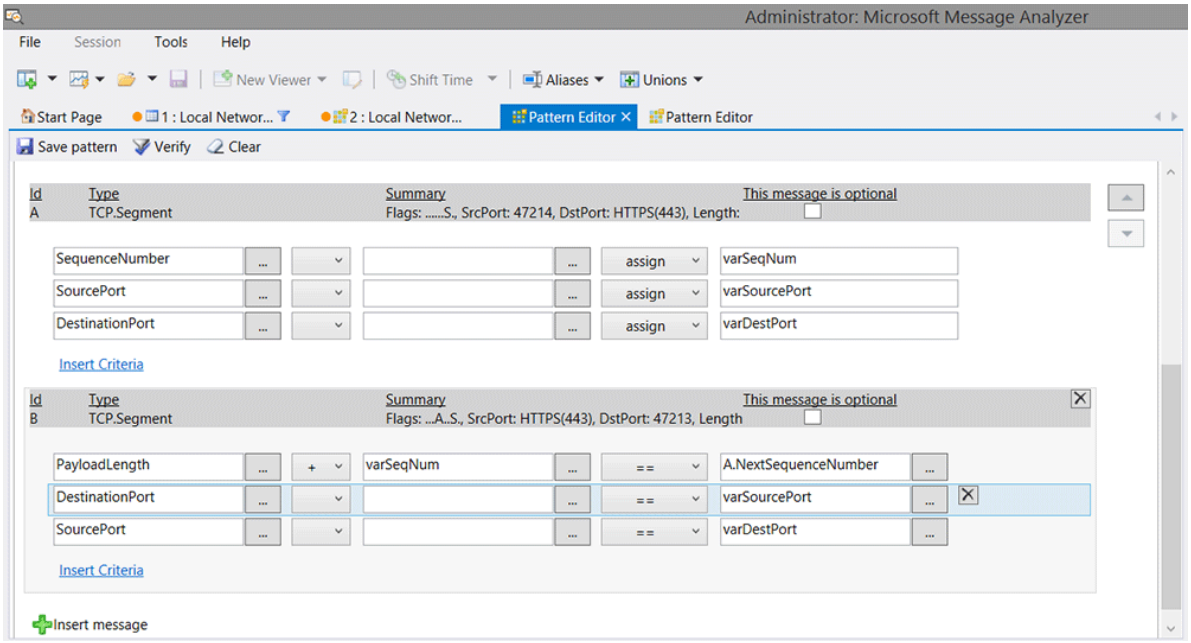

Figure 44: Example of building a Pattern expression

13. Click the **Verify** button on the **Pattern Editor** toolbar and confirm that you receive no compilation errors.
14. Specify a name for this **Pattern** expression in the **Name** text box of the **Pattern Editor**, for example, "TCP SeqNum Pairs".
15. Optionally, provide a description for this **Pattern** expression in the **Description** text box.
16. Click the **SUMMARY** subtab on the **Quick** tab of the **Pattern Editor** dialog and place a check mark in each check box to add the indicated values as columns of data for the **MATCHED INSTANCES** pane of the **Pattern Match** viewer.
17. Click the **Save pattern** button on the toolbar of the **Pattern Editor** dialog to save this **Pattern** expression.

> **NOTE**
> As soon as you save a new **Pattern** expression, Message Analyzer automatically executes the **Pattern** expression code against the current set of trace results.

If you want to view the **Pattern** expression code that Message Analyzer created in the background for the criteria control fields, operators, and values that you specified, you can hover over the **Pattern** expression in the **AVAILABLE PATTERNS** list of the **Pattern Match** viewer with your mouse, or you can right-click the **Pattern** expression and select the **Edit** or **Create a Copy** command in the context menu that appears, to display the code on the **Free Form** tab of the **Pattern Editor** dialog.

> **TIP**
> After you save a **Pattern** expression, you can leave the **Pattern Editor** dialog open so you can revisit the prepopulated view, should you need to alter the configuration with the automation controls. If you close the **Pattern Editor** for a newly created **Pattern** expression, you can then only view the **Pattern** expression in the **Free Form** configuration mode, which displays the OPN code rather than the automation controls that are accessible on the **Quick** tab .

Editing Pattern Expressions

You can edit, delete, set as a favorite, or make a copy of any **Pattern** expression that appears in the **My Items** category of your local **Pattern** expression Library in the **AVAILABLE PATTERNS** pane of the **Pattern Match** viewer. However, for the predefined **Pattern** expressions that appear in the **Network** category, you can only make a copy of these since they cannot be modified or deleted. To edit a **Pattern** expression, you will use the same **Pattern Editor** already described herein.

Editing a User Defined Pattern Expression
To start editing a **Pattern** expression in the **My Items** category, right-click it in the **AVAILABLE PATTERNS** list in the **Pattern Match** viewer. This action launches the **Pattern Editor** with the **Free Form** tab open and the **Quick** tab disabled. From here, you can change the **Name**, **Description**, and **Category**; modify the OPN code in *free form* style; and open the **Pattern** expression Library to set the **Favorites** status of an expression. You can also **Delete** any **Pattern** expression that you previously created, or you can **Create a Copy** of it, for example, to use as a template for another **Pattern** expression that you want to create.

Creating a Copy of a Built-In Pattern Expression
For any of the built-in **Pattern** expressions that exist in the **Network** category of the **AVAILABLE PATTERNS** list, you can use the right-click **Create a Copy** command to utilize the code of such an expression, or a portion thereof, as a template for another **Pattern** expression that you want to create. When you select the **Create a Copy** command from the context menu that appears when you right-click a built-in **Pattern** expression, the **Pattern Editor** displays with the **Free Form** tab open and the **Quick** tab disabled. From here, you can make a copy of the expression, modify the code, change the **Name**, **Description**, and **Category**, set its **Favorites** status, and **Save** it as a new **Pattern** expression that becomes part of **Message Analyzer Sequence Expressions** asset collection Library that displays in the **AVAILABLE PATTERNS** list.

Saving a Pattern Expression

When you finish configuring your **Pattern** expression, click the **Verify** button on the **Pattern Editor** toolbar to ensure that you have a valid configuration, or you will be unable to save it. After you **Save** the new **Pattern** expression, OPN code is generated based on your input parameters. At that time, your **Pattern** expression is added to the **My Items** category of the **AVAILABLE PATTERNS** list that displays in the **Pattern Match** viewer. Because a saved **Pattern** expression becomes part of your local **Pattern** expression Library, you can take advantage of the management operations associated with the Message Analyzer sharing infrastructure to share your **Pattern** expressions with others, as described in Managing Pattern Expressions.

More Information
To learn more about language requirements, constructs, and other details you need to create OPN behavioral scenarios and virtual operations for pattern matching with **Pattern** expressions, see the OPN Programming Guide.

Managing Pattern Expressions

2 minutes to read

Message Analyzer enables you to save OPN behavior scenarios as **Pattern** expression items that you can share with others. Whenever you create a new **Pattern** expression or edit a copy of an existing one, you can save it to the local **Pattern** expression Library. Message Analyzer enables you to share your **Pattern** expression items with others by providing you with the capability to manage this Library. The management capabilities are provided in the **Manage Pattern** dialog that contains similar functionality that you use to manage any of the Message Analyzer asset collections that are described in Saving Settings.

Managing and Sharing Pattern Expressions

To manage your **Pattern** expression items in the **Message Analyzer Sequence Expressions** asset collection Library, click the **Manage Patterns** button on the toolbar of the **Pattern Match** viewer. This action will display the **Manage Pattern Expression** dialog, from where you can export and import selected **Pattern** expressions. You can also **Delete** any **Pattern** expression in the **My Items** category from the **Manage Pattern Expression** dialog. You can export **Pattern** expressions directly to a file share to provide other users with access to your expressions, or you can import **Pattern** expressions directly from a file share to obtain access to the expressions of other users.

You can also export **Pattern** expression items through a user feed that you configure from the **Settings** tab of the **Asset Manager** dialog, in which you point your feed to a file share or other designated location where you will export your items; the **Asset Manager** is accessible from the global Message Analyzer **Tools** menu. Thereafter, your **Pattern** expression items are accessible to any Message Analyzer user that can access the feed. Your feed will appear on the **Downloads** tab of the Message Analyzer **Asset Manager** dialog, where other users can access a consistent interface that enables them to download your postings. However, if you want to enable users to synchronize with updates to your **Pattern** expression items, some manual configuration is required, as described in Manual Item Update Synchronization.

In addition, Microsoft provides a default **Message Analyzer** feed on the **Downloads** tab of the **Asset Manager** dialog that enables you to download the **Message Analyzer Sequence Expression** asset collection from a Microsoft web service and to synchronize with collection updates that are periodically pushed out by the service. At any time, you can perform a download of an auto-synced collection from the **Settings** tab on the **Asset Manager** dialog. Thereafter, you can access and apply the downloaded or synchronized items from the local **Pattern** expression Library in the **AVAILABLE PATTERNS** list of the **Pattern Match** viewer.

More Information

To learn more about the common **Manage *<AssetType>*** dialog and the management capabilities it provides for Message Analyzer Library asset collections, see Managing User Libraries.

To learn more about the Message Analyzer Sharing Infrastructure, see Sharing Infrastructure, Creating Custom User Feeds, and Sharing Asset Collections on a User File Share.

Gantt Viewer

4 minutes to read

Message Analyzer provides the **Gantt** viewer, which is currently a preview feature. If you wish to use the **Gantt** viewer, you will need to select its check box on the **Features** tab of the **Options** dialog, which is accessible from the global **Tools** menu, and then restart Message Analyzer. It will then be available for selection in the **New Viewer** drop-down list on the global Message Analyzer toolbar.

Understanding the Gantt Viewer

The **Gantt** viewer is an alternate graphic data visualizer that you can select when analyzing your message data. The **Gantt** viewer provides an at-a-glance view of message dispersion across a trace timeline that is presented as color-coded protocol module identifiers, with source/destination address message pairs in the y-axis orientation and timestamps in the x-axis orientation. The **Gantt** viewer also provides axis expansion and contraction controls to facilitate zooming into and out of data points. The **Gantt** viewer enables you to determine how long any operation takes to complete with respect to **Timestamp** axis values in the display and provides this data in the context of the completion times of other operations. The figure that follows shows an example of the **Gantt** viewer.

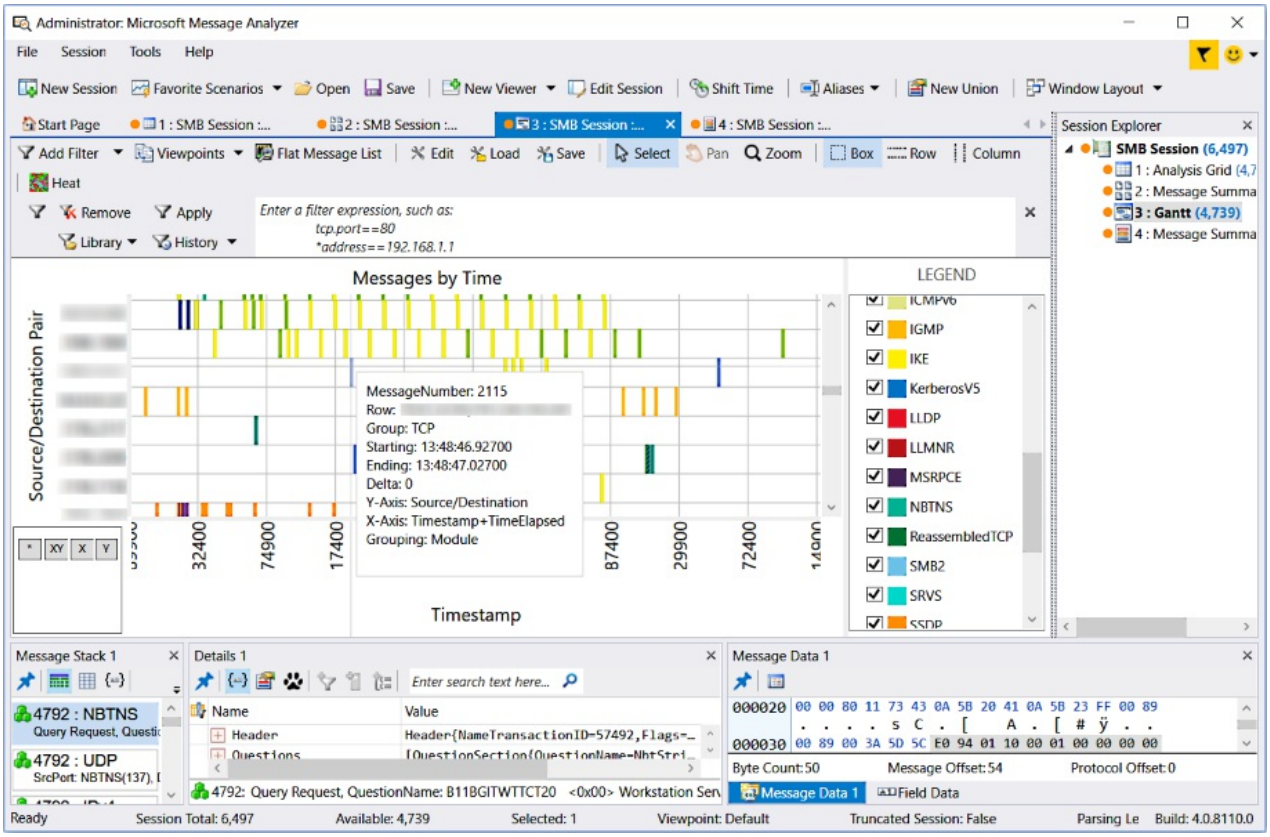

Figure 45: Gantt Viewer

Using the Toolbar Commands

To enhance data analysis capabilities the **Gantt** viewer toolbar provides the following features:

- **Edit** — click this button to open the **Gantt Configuration Editor** dialog, from where you can reconfigure the viewer layout, which includes the **Horizontal Axis** and **Vertical Axis Label** and **Layout** configuration; and the **Legend** and **Heatmap** configurations.

- **Load** — enables you to load data into the current **Gantt** viewer display from a previously saved .gantt file.
- **Save** — enables you to save data from the current **Gantt** viewer display to a .gantt file.
- **Select** mode — click this button to select single messages in the Gantt analysis surface, for example, to display message field information in the **Details Tool Window**. The **Select** mode works together with the **Box**, **Row**, and **Column** toolbar items, which you can click individually to alter the highlighting orientation of the selection feature, as described in Using the Context Menu Commands.
- **Pan** — not used.
- **Zoom** mode — after clicking this button, you can zoom into one or more displayed messages by dragging a selection box around them, for example, to select a message and view **Details** and other data.
- **Box**, **Row**, and **Column** options — see the descriptions under **Selection Mode** in Using the Context Menu Commands.
- **Heat** — click this button to display the **Heatmap Configuration** dialog, from where you can specify the following settings:
  - **Use Heatmap** — when you select this check box in the **Heatmap Configuration** dialog, the settings you make in the dialog take effect after you click **OK** to exit the dialog.
  - **Field** — by clicking the ellipsis (**...**) to the right of the **Field** text box, you can display the **Field Chooser Tool Window**, from where you can select an annotation, property, flag, or other data field for any message in the dialog that you want to highlight with **Heatmap** settings in the analysis surface.
  - **Observed** — provides an indication of the number of messages highlighted with the specified **Heatmap** settings.
  - **Minimum/Maximum Heat** — provides writeable text boxes in which you can specify **Minimum** and **Maximum** values for the **Heatmap** highlight settings.

Using the Context Menu Commands

The **Gantt** viewer also provides several context menu commands that are available by right-clicking anywhere on the main analysis surface. The commands that display and the functions they perform are described as follows:

- **Mouse Mode** — provides a submenu that contains the **Select** and **Zoom** items, where one of these modes will always be selected.
- **Selection Mode** — provides a submenu that has the following options that you can use as indicated:
  - **Box** — sets the selection mode to the box shape for selecting one or more messages nodes.
  - **Row** — sets the selection mode to enable you to select messages in a **Source/Destination Pair** row orientation.
  - **Column** — sets the section mode to enable you to select messages in a **Timestamp** column orientation.
- **Zoom** — provides several **Zoom** presets that enable you to alter the data presentation to focus on one or more messages.
- **Settings** — enables you to change the data presentation format by alternately enabling or disabling various chart elements from submenu items that include **Fixed Grid**, **Auto Scroll**, and **Display Limit** settings.
- **View Message Range in 'Analysis Grid'** — opens a selected range of messages in a separate instance of

the **Analysis Grid** viewer.

- **View Time Range in 'Analysis Grid'** — opens selected messages that fall within **Gantt Timestamp** column lines in a separate instance of the **Analysis Grid** viewer.
- **Add Rows to Filter** — enables you to create a view **Filter** from selected messages in particular rows.
- **Opaque Items** — enables you to alter the data presentation format by alternately applying and removing an opaque value to all message colors in the **Gantt** chart grid.

Legend

The **Gantt** viewer also contains a **Legend** that provides a list of check boxes that correspond to the protocols or modules for which messages were captured in the trace results displayed in the **Gantt** analysis surface. To focus on messages from one or more particular protocols or modules, unselect all the ones for which you do not want to view data.

Chart Viewer Layouts

5 minutes to read

By default, Message Analyzer provides several built-in data viewers such as the **Analysis Grid**, **Grouping**, **Pattern Match**, and **Gantt** viewers, as described in earlier sections. However, Microsoft also provides numerous view **Layouts** for the **Chart** viewer that each have a predefined configuration designed to display specific types of data in various presentation formats for the enhancement of data analysis perspectives. The built-in **Layouts** for the **Chart** viewer consist of various graphic visualizer components that provide top-level data summaries and statistics based on message types, fields, properties, and preconfigured formulas. The visualizer component types include **Table** grid, **Bar** element, **Pie** chart, **Timeline** graph, and other graphs that plot message data in the x-y coordinate domain; the **Timeline** graphs also include time window slider controls for zooming on data. A **Chart** viewer **Layout** called the **Top Level Protocols Message Count** contains a **Bar** element visualizer component, as shown in the figure that follows. The data exposed in this **Layout** shows the relative distribution of message volume for all protocols or modules for which messages were captured in a set of trace results.

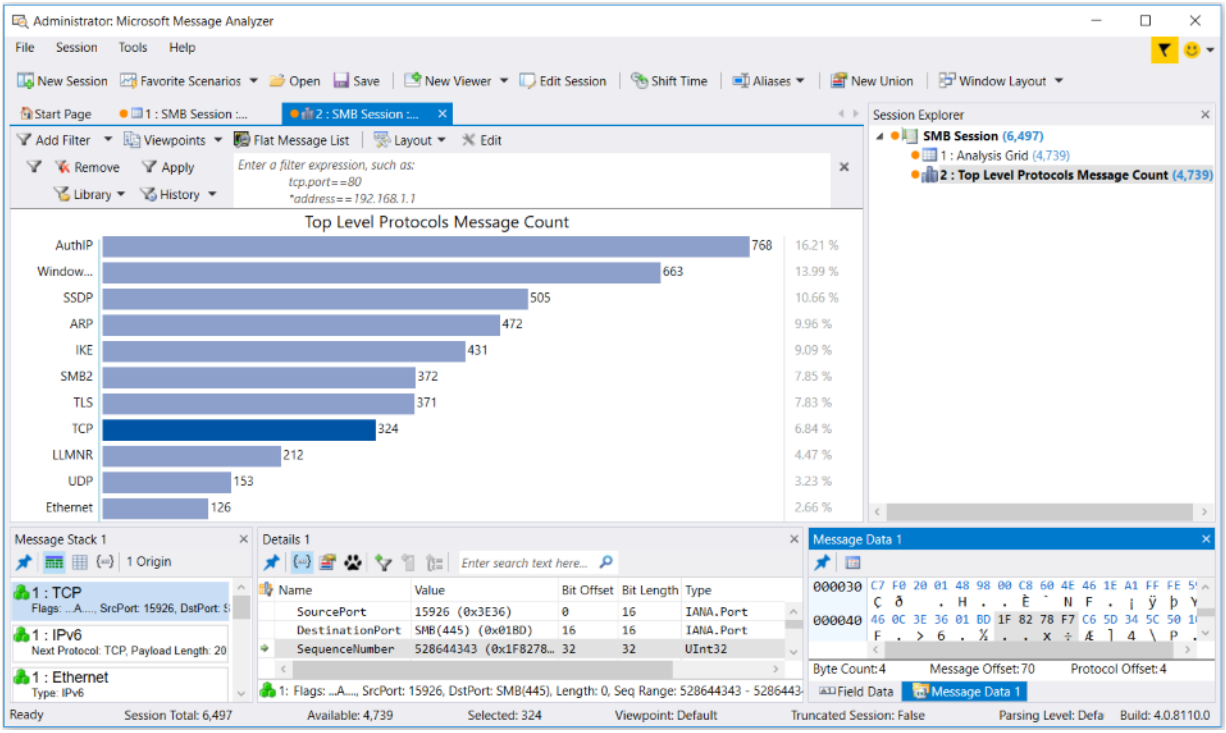

Figure 46: Top Level Protocols Message Count Chart Layout

The **Layouts** feature for **Charts** enables you to select from a wide assortment of built-in view **Layouts** that are custom designed by Microsoft to provide focused analysis environments for quick exposure of information that is critical to troubleshooting. Many of the **Layouts** provide high-level overviews of data in unique formats that can immediately point to specific issues and potentially indicate the direction in which further analysis might proceed. The categories in which built-in **Chart** viewer **Layouts** exist are described in the Chart Layout Categories section ahead, where you can link to topics that describe the **Chart** viewer **Layouts** in each category.

A single type of graphic visualizer now exists in each built-in **Layout**, such as a grid, bar element, pie-chart, or timeline component. The data points that you will find in many of these components are dynamically integrated with other viewers such as the **Analysis Grid**, so that you can interactively drive the display of messages in the latter by double-clicking data points in the former. You can also click a group node in the **Grouping** viewer to simultaneously drive the display of messages in the **Analysis Grid** viewer, a **Chart**, and

in a **Tool Window** such as the **Message Stack**. In any individual view of data, you can achieve a unique analysis context; however, you will find that a combination of **Chart Layouts**, viewers, and **Tool Windows** that expose the same session data or details in different ways will contribute significantly to problem identification and resolution.

Chart Viewer Layouts in Message Analyzer Profiles
In the case of **Profiles**, the **Chart** viewer **Layouts** are preselected to coordinate with the **Layouts** of other viewers, such as the **Analysis Grid** and **Grouping** viewers. These preset viewer and layout configurations are carefully chosen by Microsoft for each **Profile**, so that you can achieve a unified analytical environment after you load data into Message Analyzer while a particular **Profile** is enabled. Applied **Profiles** create these environments with the use of view **Layouts** that are data complementary and interactive data viewers that enable you to correlate data across different presentation contexts. With the preset data viewer and layout configurations of a **Profile**, you can obtain a multi-faceted analysis perspective based on data displays that are rich with information, that is, after you load data into Message Analyzer from the file type with which the **Profile** is associated.

More Information
To learn more about **Profiles** and how you can use some of the viewer **Layouts** that they employ for analysis and troubleshooting purposes, see Working With Message Analyzer Profiles.

Creating Custom Chart Viewer Layouts

Message Analyzer also enables you to create custom **Layout** configurations by editing any built-in **Chart** viewer **Layout** and saving it under a different name. When creating a new **Layout**, you can choose which visualizer component you want to use by selecting a **Bar**, **Pie**, **Timeline**, or **Grid** style visualizer component. You can also choose the fields and formulas that provide the **Layout** configuration and functions that you want. However, to configure settings, you must have an in-focus **Chart** viewer **Layout** already displayed in order to get access to the **Edit** command that appears in the **Chart** drop-down on the global Message Analyzer **Session** menu. After you click the **Edit** command, the **Edit Chart Layout** dialog displays from where you can create and **Apply** your custom configuration. Thereafter, you can save the **Layout** with the use of the **Save Current Layout As** command.

More Information
To learn more about configuring and saving a custom **Chart** viewer **Layout**, see Configuring Chart Viewer Layouts.

Locating the Built-In Chart Viewer Layouts

The built-in **Chart** viewer **Layouts** that are provided with every Message Analyzer installation, whether you newly install or upgrade Message Analyzer, are accessible from the following locations:

- **Start With** drop-down list — in the **New Session** dialog, you can select a default **Chart** item from the **Start With** drop-down list that displays a top-level protocols message count **Layout**, which uses a bar element configuration.
- **Layout** drop-down list — from the **Chart** drop-down list in the Message Analyzer global **Session** menu, you can access the **Layout** drop-down list from where you can select a chosen **Layout**. However, you must have a **Chart** viewer **Layout** already displayed and in-focus to access the **Layout** drop-down list. Note that this list contains different categories in which the **Layouts** are specified, as described in the Chart Layout Categories section.
- **New Viewer** drop-down list — from the **New Viewer** drop-down list on the global Message Analyzer toolbar, you can access the **Chart** drop-down list that contains all the **Layouts** that are available to

choose.

- **Session Explorer** context menu — by right-clicking anywhere in the **Session Explorer Tool Window**, you can select the **New Viewer** item in the context menu that appears to access the **Chart** drop-down list.

Obtaining New or Updated Chart Viewer Layouts

Microsoft may occasionally provide updates to the **Message Analyzer Chart View Layouts** asset collection that you can access from the **Asset Manager** dialog in the global Message Analyzer **Tools** menu; note that this dialog exposes the Message Analyzer Sharing Infrastructure. Updates to this asset collection can include new or revised **Chart** viewer **Layout** configurations that are developed at Microsoft. If you auto-sync to this collection in the **Asset Manager** dialog, you will periodically receive automatic updates and downloads of the **Message Analyzer Chart View Layouts** asset collection into your local **Chart Layouts** Library.

Chart Layout Categories

The built-in **Chart** viewer **Layouts** are described in the following sections. These section titles also reflect the actual category names in which the layouts exist in the **Layouts** drop-down that is accessible from the global Message Analyzer **Session** menu.

HTTP Category
General Category
Network Category
Netlogon Category
Networking Category
Common Category
File Sharing Category

More Information
To learn more about editing, creating, and managing **Chart** viewer **Layouts**, see Extending Message Analyzer Data Viewing Capabilities.
To learn more about how to receive downloads and updates for asset collections, see Managing Asset Collection Downloads and Updates.

See Also

Working With Message Analyzer Profiles

HTTP Category

2 minutes to read

The built-in viewer **Layouts** for **Charts** that are provided with Message Analyzer are accessible from the locations specified in the Chart Viewer Layouts topic. The viewer **Layouts** for **Charts** that are included in the **HTTP** category of the **Message Analyzer Chart View Layouts** asset collection Library consist of the following, as described in this section:

HTTP Content Type Payloads
HTTP Content Type Volumes

HTTP Content Type Payloads

2 minutes to read

The **HTTP Content Type Payloads** viewer **Layout** for **Charts** displays payloads in bytes and message count data that is associated with HTTP content types detected in a set of trace results. From this visualizer component, you can obtain the following data for analysis:

- The volumes of HTTP content types in terms of payload lengths in bytes for each type.
- The count of messages associated with each content type.

The **HTTP Content Type Payloads** data is displayed in tabular format and includes **ContentType**, **PayloadBytes**, and **MessageCount** columns that contain the data. Note that the bar element graph in the **HTTP Content Type Volumes** viewer **Layout** repeats some of this data but does not include the **MessageCount** for each content type.

Analyzing Payload Data

The tabular data of this **Chart** can help you see at a glance which payload byte values and message counts are the largest for any particular content type. In turn, this can provide an indication of the loads being carried by responding web servers. For example, you might double-click a table row with the highest payload byte value to display the associated messages in a separate instance of the **Analysis Grid** viewer. You can then right-click the **Destination** column in the **Analysis Grid** viewer and select the **Group** command to create groups of different web server names along with the number of messages associated with each one. This can tell you immediately which server/s are carrying the highest load for a particular content type, which in turn can provide an indication of web server performance.

See Also

HTTP Content Type Volumes

HTTP Content Type Volumes

2 minutes to read

The **HTTP Content Type Volumes** viewer **Layout** for **Charts** displays **HTTP Content Type Volumes** in a horizontal bar element graph. This visualizer component provides the following for a set of trace results:

- The relative volumes of HTTP content types in terms of payload lengths in bytes for each type.
- A visual indication that shows the relative distribution of byte volume for each content type.

The labels to the left of the bar graph indicate the content type, the values to the right of the graph represent percentage values that are relative to the total volume in bytes for all payloads, and the bar elements of the graph provide a visual indication of such volumes for each content type. Note that the table in the **HTTP Content Type Payloads** viewer **Layout** repeats some of this data but also includes the total **MessageCount** for each content type.

Analyzing Volume Data

The bar element visualizer component of this **Chart** can help you see at a glance which byte volumes are the largest for any particular content type. In turn, this can provide an indication of the loads being carried by responding web servers. For example, you might double-click the bar element with the highest byte volume to display the associated messages in a separate instance of the **Analysis Grid** viewer. You can then right-click the **Destination** column in the **Analysis Grid** viewer and select the **Group** command to create groups of different web server names along with the number of messages associated with each one. This can tell you immediately which server/s are carrying the highest load for a particular content type, which in turn can provide an indication of web server performance.

See Also

HTTP Content Type Payloads

General Category

2 minutes to read

The built-in viewer **Layouts** for **Charts** that are provided with Message Analyzer are accessible from the locations specified in the Chart Viewer Layouts topic. The viewer **Layouts** for **Charts** that are included in the **General** category of the **Message Analyzer Chart View Layouts** asset collection Library consist of the following, as described in this section:

Average Elapsed Time for Operations
Average Response Time for Operations
Cluster Levels
Event Log IDs
IP/Ethernet Conversations by Message Count
IP Ethernet Conversations by Message Count Top 20
TCP/UDP Conversations by Message Count
TCP/UDP Conversations by Message Count Top 20
Top Level Protocols Message Count
Top Level Protocols Message Count Over Time

Average Elapsed Time for Operations

2 minutes to read

The **Average Elapsed Times for Operations** viewer **Layout** for **Charts** provides a horizontal **Bar** element visualizer component to display the average elapsed time for Operations that is calculated across all instances of specific methods or commands of the same type that are detected in a set of trace results, that is, for protocols that make use of request/response pairs. For example, an HTTP GET or POST method is a request message that is sent by a client to a web server, where the client awaits a response from the web server. When a response is received, Message Analyzer encapsulates the request and response messages in an Operation node that displays as a top level message row in the **Analysis Grid** viewer. The Average Elapsed Time is calculated as the average time it took for all request messages in a set of trace results to receive the *first* server response, plus the time it took to receive all message fragments associated with the server response messages.

Understanding the Average Elapsed Time

The **Average Elapsed Time** for each method or command detected in a set of trace results measures the difference between the time that a Request method was sent by the client to a particular server and the time at which the last message fragment associated with the Request was received by the client. Therefore, the elapsed time indicated by a bar element for a specific method or command provides an average value that may indicate when there are network latency issues, providing that server **Average Response Times** are clearly low values (which generally rule out slow server responses). Because the bar graph data is sorted from the longest to the shortest average elapsed times, you can quickly ascertain from this data whether network issues are impacting performance. The Average Elapsed Time for each method or command is specified on the left side of the bar graph, while on the right side there is a percentage value for each method or command that is relative to the total elapsed time for all the methods or commands in the trace. This can provide a quick indication of which methods or commands are taking the longest time to complete transmission of all related data. With this information, you might focus your troubleshooting efforts on network performance.

Note that the Average Elapsed Time values depicted in the **Average Elapsed Times for Operations** viewer **Layout** are the same as the values in the **TimeElapsed** column of the **Analysis Grid** viewer.

> **NOTE**
>
> You can double-click any bar element in either of the graphs in this viewer and display the associated messages for each method or command in a separate **Analysis Grid** viewer instance for further analysis.

See Also

Average Response Time for Operations

Average Response Time for Operations

2 minutes to read

The **Average Response Times for Operations** viewer **Layout** for **Charts** provides a horizontal **Bar** element visualizer component to display the average response time that is calculated across all instances of specific methods or commands of the same type that are detected in a set of trace results, that is, for protocols that make use of request/response pairs. For example, an HTTP GET or POST method is a request message that is sent by a client to a web server, where the client awaits a response from the web server. When a response is received, Message Analyzer encapsulates the request and response messages in an Operation node that displays as a top-level message row in the **Analysis Grid** viewer. The **Average Response Time** is calculated as the average time it took for all request messages in a set of trace results to receive the first server response.

Understanding the Average Response Time

The **ResponseTime** annotation from **Field Chooser** enables you to measure the difference between the time that a Request message was sent by the client to a particular server and the first Response message received by the client from the server. Therefore, the **Average ResponseTime** indicated by a bar element for a specific method or command provides an average value that tells you how long a particular server is taking to reply to requests. Because the bar graph data is sorted from the longest to the shortest average response times, you can quickly ascertain from the upper rows of data which requests are taking the longest time to get a response for a particular method or command, possibly pointing to server performance problems. The **Average Response Time** that is associated with each method or command is also specified on the right side of the bar graph as a percentage value that is relative to the total response time for all Operation messages in a trace. This percentage value can also provide a quick indication of which server/s are taking the longest time to respond. With this information, you can focus your troubleshooting efforts on performance of the server that is sending response messages to the client.

To determine which server may be having performance problems, you might be able to isolate that information through filtering. For example, for HTTP GET methods that are reporting a long average response time, you might apply a view **Filter** to the **Analysis Grid** viewer such as `HTTP.Method == "GET"` to isolate the Operations that contain the HTTP requests. You can then add the **ResponseTime** annotation from the **GlobalAnnotations** node in **Field Chooser** as a new column in the **Analysis Grid** viewer, from where you can correlate the high response time values with the IP address of the web server in the **Destination** column. You might perform a similar filtering operation for an `SMB2.ComNegotiate` request message to a file server.

See Also

Average Elapsed Time for Operations

Cluster Levels

3 minutes to read

The **Cluster Levels** viewer **Layout** for **Charts** provides a horizontal **Bar** element visualizer component that displays the log entry count for each of the information levels (**InfoLevel**) that exist in a Cluster.log file that were reported by components of the Cluster Service. The typical values for **InfoLevel** consist of **DBG**, **INFO**, **WARN**, and **ERR**. These **InfoLevel** values provide debug information, status information, warnings, and errors that were reported to the Cluster log by various subcomponents of the Cluster Service, such as the Global Update Manager (GUM), Failover Manager (FM), Node Manager (NM), and Database Manager (DM) services.

With the **Cluster Levels** viewer **Layout**, you can obtain a quick assessment of which information levels have the most log entry activity. The bar element configuration of this **Layout** provides an at-a-glance summary of the relative distribution of log entry volume for each of the information types found in a Cluster log file. The information level values of the most interest are likely to be the warnings (**WARN**) and errors (**ERR**), which can point you to areas in the Cluster Service on which your analysis should focus.

> **NOTE**
> A Cluster.log file will be parsed only if you select the **Cluster** configuration file in **Text Log Configuration** drop-down list of the **New Session** dialog for a Data Retrieval Session prior to loading the data into Message Analyzer. Otherwise, no data will display in the **Cluster Levels** view **Layout**.

More Information
To learn more about working with text-based .log files, see Opening Text Log Files.

Using the Cluster Levels Layout

When you review the data presented in the **Cluster Levels** viewer **Layout**, you might notice a significant volume of log entries with **ERR** levels being reported. To assess these entries, you might proceed by double-clicking the **ERR** bar element to isolate all the log entries that contain the **ERR InfoLevel** type in a new instance of the **Analysis Grid** viewer. Then you can sort the **SubComponent** column of the **Analysis Grid** viewer in ascending order to consolidate all the subcomponents of the same type together for easier viewing. Thereafter, you can review the **RemainingText** column for descriptions of the errors and failures that occurred for each subcomponent of the Cluster Service. Note that you might also execute the **Group** context menu command on the **Subcomponent** column of the **Analysis Grid** viewer to create a grouped view of the subcomponents and the associated log entry counts for an enhanced perspective of the data. Moreover, you can use the **Cluster Logs Layout** for the **Grouping** viewer to create a hierarchical grouped configuration of nested **InfoLevel**, **Subcomponent**, and **ProcessId** groups, so that you can quickly isolate the data according to **InfoLevel** values and view the underlying **Subcomponents** that logged the debugging information.

Interactive Analysis
This **Layout** for the **Chart** viewer is intended to work with the **Cluster Log Layout** for the **Analysis Grid** viewer and the **Cluster Logs Layout** for the **Grouping** viewer to create an integrated and interactive analysis environment. You will be able to correlate the data most effectively if you have these viewers and **Layouts** displayed. Note that these viewers and **Layouts** are configured in the **Cluster Logs Profile** and will display after you load data from a Cluster.log file, provided that you enabled this **Profile** on the **Profiles** tab of the **Options** dialog (accessible from the global Message Analyzer **Tools** menu).

As an example of interactively driving the display of messages, if you select an **InfoLevel** group in the **Grouping** viewer, you can display all the messages associated with that particular error level in the **Analysis Grid** viewer for

further analysis of details.

More Information

To learn more about interactively analyzing Cluster log data with the indicated viewing and **Layout** configurations, see the following topics:

Applying and Managing Analysis Grid Viewer Layouts — see the **Cluster Log Layout** in this topic.

Grouping Viewer — see the **Cluster Logs Layout** in this topic.

Working With Message Analyzer Profiles — see the **Cluster Logs Profile** in the table of this topic to review a related usage scenario and analysis example and to learn how to manually display the **Grouping** and **Chart** viewers with the **Layouts** defined in this **Profile**.

Event Log IDs

2 minutes to read

The **Event Log IDs** viewer **Layout** for **Charts** provides a horizontal **Bar** element visualizer component that displays the message count associated with each **EventID** in an event (*.evtx) log. This **Layout** provides an immediate visual assessment of the relative distribution of message volume per **EventID** — ordered from the highest to the lowest volume. As you might expect, this exposes which **EventIDs** had the most significant volumes.

Using the Event Log IDs Layout

From a quick visual assessment of the data in this **Layout** alone, you can determine the events that involved the highest message count, which could be an indication of where further investigation is needed, for example, very busy processes that are generating a large number of events. High volumes might also point to an overburdened system component or application that is issuing a lot of event traffic or experiencing a high rate of errors. Sparse traffic might be an indication of dropped packets due to misconfigured ETW Session buffer settings, as described in Specifying Advanced ETW Session Configuration Settings.

Interactive Analysis

The **Event Log IDs Layout** for the **Chart** viewer is intended to work with the **Event Log Layout** for the **Analysis Grid** viewer and the **Event Viewer Layout** for the **Grouping** viewer, to create an integrated and interactive analysis environment. You will be able to correlate the data most effectively if you have these viewers and **Layouts** displayed. Note that these viewers and **Layouts** are configured in the **Event Logs Profile** and will display after you load data from a *.evtx file, provided that you enabled this **Profile** on the **Profiles** tab of the **Options** dialog (accessible from the global Message Analyzer **Tools** menu).

As an example of interactively driving the display of messages associated with a particular **EventID**, you can double-click any bar element in the **Event Log IDs Layout** to display the messages associated with that element in a new instance of the **Analysis Grid** viewer for focused analysis of the messages associated with a particular **EventID**. The **Event Viewer Layout** for the **Grouping** viewer also enables a greater interactive context with additional enhancements to the analysis process.

More Information

To learn more about interactively analyzing Event log data with the indicated viewing and **Layout** configurations, see the following topics:

Applying and Managing Analysis Grid Viewer Layouts — see the **Event Log Layout** in this topic.

Grouping Viewer — see the **Event Viewer Layout** in this topic.

Working With Message Analyzer Profiles — see the **Event Logs Profile** in the table of this topic to review a related usage scenario and analysis example and to learn how to manually display the **Grouping** and **Chart** viewers with the **Layouts** defined in this **Profile**.

IP-Ethernet Conversations by Message Count

2 minutes to read

The **IP/Ethernet Conversations by Message Count** viewer **Layout** for **Charts** provides a summary of the top Network layer conversations that took place within the time boundaries of a set of trace results. The summary data is presented in a **Table** grid visualizer component that provides message details in tabular format and exposes useful statistics for the top Network layer conversations that took place in a set of trace results. The data is sorted in descending order from the highest to lowest message count, per conversation. The **Table** grid contains the following data columns:

- **Network** — contains the IP or Ethernet addresses of the computer nodes that exchanged messages.
- **Count** — specifies the number of messages that were exchanged in each conversation.
- **Bytes** — based on the `payloadLength` property of any module that defines this field, the values in this column specify the total payload byte volume of all messages (containing this property) that are associated with each Network conversation.
- **KBs** — provides an indication of the data transmission rate in kilobytes-per-second for the messages in a conversation.
- **Duration** — specifies the delta between the **StartTime** and **EndTime** values.
- **StartTime** — specifies the time at which the first message in a conversation group began.
- **EndTime**— specifies the time at which the last message in a conversation group ended.
- **BPS** — provides an indication of the data transmission rate in bytes-per-second for the messages in a conversation.
- **K** — a chart constant to facilitate calculation of the KBs values.

Using the IP/Ethernet Conversations by Message Count Layout

To maximize analysis effectiveness, you should apply the **Network Viewpoint** to this **Layout**, which removes all messages above the Network layer. When you apply this **Viewpoint** from the **Viewpoint** drop-down list on the Filtering toolbar, the values in columns such as **Count**, **Bytes**, **KBs**, and **BPS**, will then be based on IP and Ethernet messages only. For example, the **Count** column will then indicate the number of IP or Ethernet messages in each conversation, while the **Bytes** column will indicate the number of IP or Ethernet payload bytes exchanged in each conversation.

Statistics that you can obtain from this grid visualizer component that are useful in troubleshooting the Network layer include:

- Message volume per conversation
- Top bandwidth consumers
- Data transmission rates
- Time to complete conversations

NOTE

You can double-click any conversation row in the grid to display the messages for the associated conversation in the **Analysis Grid** viewer for further analysis.

See Also

IP Ethernet Conversations by Message Count Top 20

IP Ethernet Conversations by Message Count Top 20

2 minutes to read

The **IP/Ethernet Conversations by Message Count Top 20** viewer **Layout** for **Charts** provides a summary of the top Network layer conversations that took place within the time boundaries of a set of trace results. The summary data for this **Layout** is provided in a **Bar** element visualizer component that displays the IP addresses of communicating computers and the number of messages exchanged in each conversation for the top 20 Network layer conversations in a trace. A scale is also provided to the right of the bar element graph that indicates the relative percent volume of messages for each bar element with respect to the overall top 20 conversations message count.

Using the IP/Ethernet Conversations by Message Count Top 20 Layout

To maximize analysis effectiveness, you should apply the **Network Viewpoint** to this **Layout**, which removes all messages above the Network layer. When you apply this **Viewpoint** from the **Viewpoint** drop-down list on the Filtering toolbar, the **Layout** values will then be based on IP and Ethernet messages only. This can result in changing the number of messages associated with one or more conversations.

Statistics that you can obtain from this **Layout** that are useful in troubleshooting the Network layer include:

- Message volume per conversation
- Top bandwidth consumers
- Relative distribution of message count per conversation as a percentage of all the messages displayed in the **Layout**.

> **NOTE**
>
> You can double-click any bar element in the **Layout** to display the messages for the associated conversation in the **Analysis Grid** viewer for further analysis.

See Also

IP/Ethernet Conversations by Message Count

TCP-UDP Conversations by Message Count

3 minutes to read

The **TCP/UDP Conversations by Message Count** viewer **Layout** for **Charts** provides a summary of the top Transport Layer conversations that took place within the time boundaries of a set of trace results. The summary data is provided in a **Table** grid visualizer component to expose message details and useful statistics for the top Transport Layer conversations that took place in a set of trace results. The data is sorted in descending order from the highest to lowest message count, per conversation. The grid layout contains the following data columns:

- **Network** — contains the IPv4, IPv6, and Ethernet addresses of the computer nodes that exchanged messages in each conversation in a set of trace results.
- **Transport** — contains the TCP or UDP ports that carried the conversations.
- **Count** — specifies the number of messages that were exchanged in each conversation.
- **Bytes** — based on the `payloadLength` property of any module that defines this field, the values in this column specify the total payload byte volume of all messages (containing this property) that are associated with each conversation.
- **KBs** — provides an indication of the data transmission rate in kilobytes-per-second.
- **Duration** — specifies the delta between the **StartTime** and **EndTime** values, which calculates the duration of each conversation in a set of trace results.
- **StartTime** — specifies the time at which the first message in a conversation group began.
- **EndTime** — specifies the time at which the last message in a conversation group ended.
- **BPS** — provides an indication of the data transmission rate in bytes-per-second.
- **K** — a chart constant to facilitate calculation of the KBs values.

Using the TCP/UDP Conversations by Message Count Layout

To maximize analysis effectiveness at the Transport Layer, the **Transport Layers TCP/UDP Viewpoint** has been applied to this **Layout** by default, which removes all messages above the Transport Layer. You can observe evidence of this background configuration in the **Viewpoints** drop-down list on the Filtering Toolbar, where you will see that the **Layout ViewPoint** item is selected, that is, after you display the **TCP/UDP Conversations by Message Count Layout**. As a result of this applied **Viewpoint**, the values in columns such as **Count**, **Bytes**, **KBs**, and **BPS**, will be based on top-level TCP and UDP messages only. For example, the **Count** column will then indicate the number of top-level TCP or UDP messages in each conversation, while the **Bytes** column will indicate the number of TCP or UDP payload bytes exchanged in each conversation. Note that you can also isolate conversations by TCP or UDP messages if you apply a **TCP** or **UDP Viewpoint**, respectively.

Statistics that you can obtain from this **Layout** that are useful in troubleshooting the Transport Layer include:

- TCP message volume per conversation
- Payload volume per conversation
- Data transmission rates per conversation
- Duration per conversation

Interactive Analysis

The **TCP/UDP Conversations by Message Count** view **Layout** for **Charts** is intended to work with the **Analysis Grid** viewer and **Grouping** viewer **Layouts** that are configured in all of the **Performance Top Down** and **Network Monitor Profiles**. These **Profiles** create an interactive and integrated analysis environment for data that exists in various trace file types, for which **Profiles** are configured. For example, the **Performance Top Down Profile** for *.cap files uses the following data viewers and **Layouts** to create an integrated analysis environment.

- **Analysis Grid** viewer — uses the **Performance Top Down Layout** in this **Profile**.
- **Grouping** viewer — uses the **Process Names and Conversations Layout** in this **Profile**.
- **Chart** viewer — uses the **TCP/UDP Conversations by Message Count Layout** in this **Profile**.

You will be able to correlate the data most effectively if you have these viewers and **Layouts** displayed. For example, you can interactively drive the display of messages associated with a particular TCP conversation into a new instance of the **Analysis Grid** viewer by double-clicking any table-grid row in the **TCP/UDP Conversations by Message Count Layout**. The results enable you to achieve a focused analysis of the messages associated with a particular conversation. The **Process Names and Conversations Layout** for the **Grouping** viewer also enables a greater interactive context with additional enhancements to the analysis process.

More Information
To learn more about the integrated analysis environments that are created by the **Profiles** that use the **TCP/UDP Conversations by Message Count** view **Layout** for **Charts**, see the **Performance Top Down** and **Network Monitor Profiles** in the Working With Message Analyzer Profiles topic. Note that you can also review related usage scenarios and analysis examples for the **Profiles** in this topic.

See Also

TCP/UDP Conversations by Message Count Top 20

TCP-UDP Conversations by Message Count Top 20

2 minutes to read

The **TCP/UDP Conversations by Message Count Top 20** viewer **Layout** for **Charts** provides a summary of the top Transport Layer conversations that took place within the time boundaries of a set of trace results. The summary data is provided in **Bar** element visualizer component and label format for the top 20 Transport Layer conversations in a trace that displays the IP addresses of the communicating computers, the TCP/UDP communication ports and module ports that the messages transited in each conversation, and the number of messages exchanged in each conversation. A scale is also provided to the right of the bar elements of the **Layout** that indicates the relative percent volume of messages for each conversation bar with respect to the overall top 20 conversations message count.

Using the TCP/UDP Conversations by Message Count Top 20 Layout

To maximize analysis effectiveness, you should apply the **Transport Layers UDP/TCP Viewpoint** to the **Layout**, which removes all messages above the Transport Layer. When you apply this **Viewpoint** from the **Viewpoint** drop-down list on the Filtering toolbar, the values for message count in each conversation can change, as the data will now be based on TCP and UDP messages only. However, by applying this **Viewpoint**, you will create focused set of message data for analysis.

Statistics that you can obtain from this **Layout** that are useful in troubleshooting the Transport Layer include:

- Message volume per conversation
- Top bandwidth consumers
- Relative distribution of message count per Transport Layer conversation as a percentage of all the messages displayed in the **Layout**.

> **NOTE**
> You can double-click any bar chart element to display the messages for the associated Transport Layer conversation in the **Analysis Grid** viewer for further analysis.

See Also

TCP/UDP Conversations by Message Count

Top Level Protocols Message Count

2 minutes to read

The **Top Level Protocols Message Count** viewer **Layout** for **Charts** illustrates message data in a horizontal **Bar** element visualizer component that provides an at-a-glance view of the relative volume of the top-level messages from different source modules or protocols in a set of trace results. If you double-click any data bar element representing a message source in this **Layout**, the details for the top-level messages represented by that particular data bar element are rendered in a new instance of the **Analysis Grid** viewer for further analysis.

Using the Top Level Protocols Message Count Layout

This **Layout** can provide an immediate indication of which modules or protocols are the highest bandwidth consumers, based on the associated message count and/or the relative percent volume with respect to total volume of messages. As such, this **Layout** provides an at-a-glance view of the distribution of message volume per module in a set of trace results.

> **NOTE**
>
> By double-clicking a bar element in this **Layout**, all messages in which the module is either a top-level message or part of message origins will display in a separate **Analysis Grid** viewer tab. Because all messages are displayed in a separate viewer tab as indicated, the message count in that viewer tab will differ from the **Total** message count specified in the **Top Level Protocols Message Count** viewer **Layout**. To see only the top-level messages in the new **Analysis Grid** viewer tab, you can apply a view **Filter** that isolates top-level messages by using the backslash symbol ("\") described in Browsing Message Origins, as follows: "\<modulename>" — where <modulename> is a placeholder for the module or protocol of interest.

Interactive Analysis

The **Top Level Protocols Message Count Layout** for the **Chart** viewer is intended to work with the **ETW Layout** for the **Analysis Grid** viewer and the **ETW Guids and IDs Layout** for the **Grouping** viewer, to create an integrated and interactive analysis environment. You will be able to correlate the data most effectively if you have these viewers and **Layouts** displayed. Note that these viewers and **Layouts** are configured in the **ETW Analysis Profile** and will display after you load data from an event trace log (*.etl) file, provided that you previously enabled this **Profile** on the **Profiles** tab of the **Options** dialog (accessible from the global Message Analyzer **Tools** menu).

As an example of interactively driving the display of messages, you can double-click any bar element in the **Top Level Protocols Message Count Layout** to display the messages in a new instance of the **Analysis Grid** viewer for focused analysis of the messages associated with a particular protocol or module. The **ETW Guids and IDs Layout** for the **Grouping** viewer also enables a greater interactive context with additional enhancements to the analysis process.

More Information

To learn more about interactively analyzing Event log data with the indicated viewing and **Layout** configurations, see the following topics:

Applying and Managing Analysis Grid Viewer Layouts — see the **ETW Layout** in this topic.

Grouping Viewer — see the **ETW Guids and IDs Layout** in this topic.

Working With Message Analyzer Profiles — see the **ETW Analysis Profile** in the table of this topic to review a related usage scenario and analysis example and to learn how to manually display the **Grouping** and **Chart** viewers with the **Layouts** defined in this **Profile**.

See Also

Top Level Protocols Message Count Over Time

Top Level Protocols Message Count Over Time

2 minutes to read

The **Top Level Protocols Message Count Over Time** viewer **Layout** for **Charts** presents data in a **Timeline** visualizer component that plots data points in the X-Y coordinate domain. This component depicts message count as X-axis module lines that indicate the number of messages that were captured for a particular module across a set of trace results, with individual message nodes indicating the points in time when messages were captured. You can view the message count at any node for any module by hovering over one with your mouse and also in the legend to the right of the graphic viewer surface. In addition, the Y-axis is calibrated to provide an indication of relative message count for any node. Also, if you double-click a node or a module line, all the top-level messages and origins associated with that module will display in a separate **Analysis Grid** viewer tab for further analysis. You can also focus on the messages from a particular module by selecting them in a legend to the right of the timeline, and clearing selections of the message types you do not want to display.

Using the Top Level Protocols Message Count Over Time Layout

This **Layout** also enables you to adjust selectable time-window slider controls so that you can do the following:

- Visually assess the times at which protocol communications occurred within the time range of a trace.
- Use manual adjustments to drill down into specific time slots for a more granular view of the message activity that transpired there.
- Use **Zoom** presets to automatically create time windows that enable you to drill down into message activity in various time slots, starting from the beginning of a trace.

  These preset values create time windows that are **1s**, **5s**, **30s**, and **1m** in length. After you specify one of these presets, you can return to the full trace time boundaries by clicking the **All** preset, or you can manually expand the trace boundaries with the time window slider controls.

See Also

Top Level Protocols Message Count

Network Category

2 minutes to read

The built-in viewer **Layouts** for **Charts** that are provided with Message Analyzer are accessible from the locations specified in the Chart Viewer Layouts topic. The viewer **Layouts** for **Charts** that are included in the **Network** category of the **Message Analyzer Chart View Layouts** asset collection Library consist of the following, as described in this section:

IIS Log HTTP Traffic Volumes
IIS Log Server Bytes by Host over Time
IIS Log Top URI Bytes
IIS Log Top URIs by Time
TCP Rate and Diagnosis
TCP Stevens Graph
Top Talkers
Top Talkers Top 20

IIS Log HTTP Traffic Volumes

3 minutes to read

The **IIS Log HTTP Traffic Volumes** view **Layout** for **Charts** enables you to obtain a high-level summary view of specific data from an IIS log file that depicts the relative distribution of HTTP traffic volume in bytes, from the highest to the lowest volume, for the cumulative IIS server responses to each client query that the server received. The **Layout** uses a **Bar** element visualizer component that provides a Y-axis label next to each bar element to display the IP address of an IIS server and the uniform resource identifier (URI) query made to the server from an HTTP client. Each horizontal graphic bar element in this **Layout** shows the cumulative HTTP traffic volume in bytes for all server responses associated with a particular query made by a client browser or other query source. The volume values in this **Layout** are based on the **sc_bytes** field for server responses, the values for which you can view in the **Details Tool Window** after selecting a log entry in the **Analysis Grid** viewer.

> **NOTE**
>
> An IIS.log file will be parsed only if you select the **IIS** configuration file in the **Text Log Configuration** drop-down list of the **New Session** dialog for a Data Retrieval Session prior to loading the data into Message Analyzer. Otherwise, no data will display in the **IIS Log HTTP Traffic Volumes** view **Layout**.

More Information
To learn more about working with text-based .log files, see Opening Text Log Files.

Using the IIS Log HTTP Traffic Volumes Layout

This **Layout** provides a quick summary of the total server response volumes in bytes that are associated with client queries requesting access to IIS server resources and services. It enables you to see at-a-glance which client queries are driving high byte volume responses from an IIS server, which may point to areas that need further investigation. For example, very high byte volume server responses could be an indication of the potential overload of an IIS server, especially if a higher than expected number of server responses are needed to service client queries.

Interactive Analysis
This **Layout** for the **Chart** viewer is intended to work with the **IIS Layout** for the **Analysis Grid** viewer and the **IIS Layout** for the **Grouping** viewer to create an integrated and interactive analysis environment. You will be able to correlate the data most effectively if you have these viewers and **Layouts** displayed. Note that these viewers and **Layouts** are configured in the **IIS Logs Profile** and will display after you load data from an IIS.log file, provided that you enabled this **Profile** on the **Profiles** tab of the **Options** dialog (accessible from the global Message Analyzer **Tools** menu).

As an example of interactively driving the display of messages, if you select any **c_ip** group in the **Grouping** viewer, you can display all the messages associated with a particular client IP address in the **Analysis Grid** viewer for further analysis and message **Details**. In addition, each Group selection that you make in the **Grouping** viewer filters the display of messages in the **IIS Log HTTP Traffic Volumes** view **Layout** for **Charts** to create a focused context for analysis.

More Information
To learn more about interactively analyzing IIS log data with the indicated viewing and **Layout** configurations, see the following topics:
Applying and Managing Analysis Grid Viewer Layouts — see the **IIS Layout** in this topic.

Grouping Viewer — see the **IIS Layout** in this topic.
Working With Message Analyzer Profiles — see the **IIS Logs Profile** in the table of this topic to review a related usage scenario and analysis example and to learn how to manually display the **Grouping** and **Chart** viewers with the **Layouts** defined in this **Profile**.

IIS Log Server Bytes by Host over Time

2 minutes to read

The **IIS Log Server Bytes by Host over Time** view **Layout** for **Charts** enables you to obtain a high-level view of specific data from an IIS log consisting of HTTP traffic volume in bytes over time. This **Layout** uses a **Timeline** visualizer component that plots byte volume data points in time, with reference to the Y-axis calibrated in bytes and the X-axis calibrated in time. The **Layout** includes a time slider control in the lower section of the **Layout** that enables you to zoom into a configured time window to create a focused analysis context. Several **Zoom** presets are also included, such as **30s**, **1m**, **30m**, and **All**.

Using the IIS Log Server Bytes by Host over Time Layout

The data that displays in this **Layout** can give you a quick summary of the points in time where HTTP traffic volumes are the highest and how those volumes varied over the timeline. You can also perform the following actions to display the indicated data:

- **Host or site name** — mouse-hover over a data line in the x-y coordinate domain of the graph to display the name of the host or site for which the data was collected.
- **Data points** — single-click a line of data to expose the data points that exist at different times.
- **Log entries** — double-click a data line to display all related log entries in the **Analysis Grid** viewer.

Along with the **IIS Log Server Bytes by Host over Time** view **Layout**, you might also consider using the **Grouping** viewer with the **IIS Layout** and the **Analysis Grid** viewer with the **IIS Layout** to achieve an integrated and interactive analysis context that can significantly enhance your analysis process.

See Also

Applying and Managing Analysis Grid Viewer Layouts
Grouping Viewer
Working With Message Analyzer Profiles

IIS Log Top URI Bytes

2 minutes to read

The **IIS Log Top URI Bytes** view **Layout** for **Charts** enables you to obtain a high-level summary view of specific data from an IIS log file that depicts the relative distribution of cumulative IIS server response volume in bytes, from the highest to the lowest volume, for each client query requesting resources or services from an IIS server. The **Layout** uses a **Bar** element visualizer component that provides a Y-axis label next to each bar element to display the uniform resource identifier (URI) query made by an HTTP client to the server. Each horizontal graphic bar element in this **Layout** shows the cumulative volume in bytes for all the server responses associated with a particular query made by a client browser or other query source. The volume values in this **Layout** are based on the **sc_bytes** field for server responses, the values for which you can view in the **Details Tool Window** after selecting an IIS log entry in the **Analysis Grid** viewer.

Using the IIS Log Top URI Bytes Layout

This **Layout** provides a quick summary of the total server response volumes in bytes for each client query that requests access to IIS server resources and services. It enables you to see at-a-glance which client queries are driving high byte volume responses from an IIS server, which may point to areas that need further investigation. For example, very high byte volume server responses could be an indication of the potential overload of an IIS server, especially if a higher than expected number of server responses are needed to service client queries.

Along with the **IIS Log Top URI Bytes** view **Layout** for **Charts**, you might also consider using the **Grouping** viewer with the **IIS Layout** and the **Analysis Grid** viewer with the **IIS Layout** to achieve an integrated and interactive analysis context that can significantly enhance your analysis process.

See Also

Applying and Managing Analysis Grid Viewer Layouts
Grouping Viewer
Working With Message Analyzer Profiles

IIS Log Top URIs by Time

2 minutes to read

The **IIS Log Top URIs by Time** view **Layout** for **Charts** enables you to obtain a high-level summary view of specific data from an IIS log file that depicts the relative distribution of the average time in milliseconds (ms), from the highest to the lowest average time, that an IIS server took to service each client query requesting resources from it. The **Layout** uses a **Bar** element visualizer component that provides a Y-axis label next to each bar element to display the uniform resource identifier (URI) query made by an HTTP client to the server. Each horizontal graphic bar element in this **Layout** shows the average service time (ms) of all the server responses associated with a particular query made by a client browser or other query source. The time values in this **Layout** are based on the **time_taken** field for server responses, the values for which you can view in the **Details Tool Window** after selecting an IIS log entry in the **Analysis Grid** viewer.

Using the IIS Log Top URIs by Time Layout

This **Layout** provides a quick summary of the average time it took for an IIS server to service each client query that requested access to its resources and services. It enables you to see at-a-glance which client queries took the longest to be serviced, which may point to areas that need further investigation. For example, very high average service times could be an indication of the potential overload of an IIS server, or possibly network latency issues when servicing client queries.

Along with the **IIS Log Top URIs by Time** view **Layout** for **Charts**, you might also consider using the **Grouping** viewer with the **IIS Layout** and the **Analysis Grid** viewer with the **IIS Layout** to achieve an integrated and interactive analysis context that can significantly enhance your analysis process.

See Also

Applying and Managing Analysis Grid Viewer Layouts
Grouping Viewer
Working With Message Analyzer Profiles

TCP Rate and Diagnosis

2 minutes to read

The **TCP Rate and Diagnoses** view **Layout** for **Charts** enables you to obtain a focused view of specific TCP data from a set of trace results that tells you the ratio of **Diagnoses** messages to all messages in each IP conversation in which **Diagnoses** errors occurred. This **Layout** uses a **Table** grid visualizer component that is driven by a data formula that automatically calculates a **Rate** value that exposes this message ratio. This **Layout** also has a **TCP Viewpoint** applied in the background, which isolates TCP messages to create a focused analysis context.

Using the TCP Rate and Diagnoses Layout

Because Message Analyzer automatically calculates the indicated **Rate** value, you can quickly narrow down the IP conversations where the highest **Rate** values exist, which in turn can provide a focus point for further investigation of TCP issues. For example, if you have a **Rate** value of .50, this means that half of the messages in a particular IP conversation contain underlying TCP **Diagnoses** messages. A **Rate** at this level could be significant, especially if the conversation also contains a high message count. To proceed with analysis, you might double-click the table grid row containing the **Rate** of interest to drive the display of conversation messages into a new instance of the **Analysis Grid** viewer for further analysis of the TCP **Diagnoses** messages and other message **Details**. From the **Analysis Grid** viewer, you can click a **Diagnoses** message icon in the **DiagnosisTypes** column to display the inline details for the **Diagnoses** message. In addition, given that the **Diagnostics Tool Window** data is driven by the content of the current in-focus data viewer, you can conveniently review the details of all **Diagnoses** messages in the IP conversation with which you are working from this **Tool Window**.

> **NOTE**
>
> If the **Diagnostics Tool Window** is not displayed, you will need to enable it on the **Features** tab of the **Options** dialog in the global Message Analyzer **Tools** menu, and then restart Message Analyzer. Thereafter, you will need to select the **Diagnostics** window from the **Windows** submenu of the global **Tools** menu to launch it.

TCP Stevens Graph

4 minutes to read

The built-in **TCP Stevens Graph** view **Layout** for **Charts** provides a graph of TCP sequence numbers versus time, along with a **Source Diagnoses** statistics table, which together can reveal certain patterns that indicate specific TCP traffic issues. The data provided by this chart can help you determine whether or not TCP traffic is transiting freely without interruptions, significant delays, or loss of packets.

The graphic portion of the **TCP Stevens Graph** view **Layout** consists of IP conversation lines of connected message nodes that correlate with y-axis sequence numbers, while the x-axis is calibrated to show message timestamps. You can click a conversation line to view the dispersion of message nodes and then hover over any message node with your mouse to see details such as the endpoints that participated in the IP conversation, the message module, source and destination ports, the timestamp, and a sequence number suffix in the current series. You can also zoom into the conversation lines to obtain greater resolution for analysis purposes. The standard zoom presets that come with Message Analyzer timeline visualizers are included. Lastly, you can hide or display conversation lines in the **TCP Stevens Graph** by unselecting or selecting the check boxes, respectively, in the legend to the right of the **TCP Stevens Graph**.

Understanding the Source Diagnoses Information

The **Source Diagnoses** table provides supplementary information that enhances your graphic chart analysis by providing you with statistics that show the IPv4 and IPv6 conversations that took place, the associated port numbers, the message count in each conversation segment, the number of diagnosis messages that occurred in each conversation segment, and a decimal value that expresses the ratio of the number of diagnosis messages to the total number of messages in each conversation segment. Note that the Message Analyzer **Grouping** feature is applied behind the scenes in this chart to create the IP conversation groups. This data is presented in tabular form by the columns of information indicated below:

> **NOTE**
> Observe that the graphic portion of the **TCP Stevens Graph** view **Layout** shows IP conversation *segments*, for example: a request/response pair can display as two separate segments; the graph legend correlates the conversation groups; and the **Source Diagnoses** table breaks the conversation groups into conversation segments for analysis.

- **Source** — specifies the IP addresses of source computers that participated in IP conversations.
- **Destination** — specifies the IP addresses of destination computers that participated in IP conversations.
- **SourcePort** — specifies the computer source ports that carried messages in an IP conversation.
- **DestinationPort** — specifies the computer destination ports that carried messages in an IP conversation.
- **Rate** — specifies the ratio of the number of diagnosis messages to the total message count for an IP conversation segment. This value measures diagnosis error messages as a fraction of the total message count. A high **Rate** value can be an indication of TCP issues such as retransmits, duplicate ACKs, lost segments, dropped packets, application glitches, and so on.
- **Count** — indicates the total number of messages in each IP conversation segment. This value is also reflected in the **Stevens** graph as the number of message nodes that display when you click a IP conversation segment line in the graph.
- **Diagnoses** — indicates the total number of diagnosis messages in each IP conversation segment. The

higher this number, the more TCP issues are present in a particular conversation.

Using the TCP Stevens Graph Layout

As indicated at the bottom of this chart, you should apply the **TCP Viewpoint** to this **Layout** to isolate TCP messages at top level with no messages above it, so that you can analyze your trace data from the perspective of TCP. To apply the **TCP Viewpoint**, click the **Viewpoint** drop-down list on the Filtering toolbar and then select the **TCP** item in the list. By applying this **Viewpoint**, the **TCP Stevens Graph** will be redrawn with data that is filtered by the **TCP Viewpoint** and the **Source Diagnoses** table will be appropriately repopulated with the filtered data. You can then begin analysis of values for signs of TCP traffic issues. Some typical patterns that you might encounter with the use of this **TCP Stevens Graph** and the issues they can expose consist of the following:

- Long gaps between message sequence numbers can indicate times when TCP message throughput is low or nonexistent.
- Long intervals between message nodes/sequence numbers can indicate transmission delays.
- Conversation segment lines that have a more vertical orientation with message nodes in close proximity on the lines indicate traffic is moving quickly. This is the ideal case.
- Conversation segment lines with a more horizontal orientation where message nodes are not in close proximity indicate it is taking more time for messages to transit. This is the problem case.
- Because the **Source Diagnoses** table indicates traffic in two directions, you can correlate the associated source and destination traffic with conversation segment lines in the **TCP Stevens Graph**, to determine whether traffic is moving faster or slower in one direction or the other.

> **NOTE**
>
> You can double-click a **Source Diagnoses** table row to display the associated messages in a separate **Analysis Grid** viewer tab.

Top Talkers

2 minutes to read

The built-in **Top Talkers** view **Layout** for **Chart** provides a summary of the endpoint conversation address pairs that had the top message volume and payload bytes values in a set of trace results. This **Layout** also displays additional statistics that provide an indication of the data transmission rate and the distribution of network busyness with respect to traffic volume and duration for each endpoint address pair. The summary data is contained in a **Table** grid visualizer component that provides endpoint address pairs, message and byte count, and other traffic details in a column layout that is tailored to provide statistics that are useful for analysis. The grid data is sorted in descending order from the highest to lowest message count and by highest to lowest payload byte count. The **Layout** contains the following data columns:

- **AddressPair** — values in this column specify the endpoint address set for which traffic statistics are generated.
- **Count** — values in this column specify the total volume of messages for each endpoint address.
- **Bytes** — based on the `payloadLength` property of any module that defines this field, the values in this column specify the total payload byte volume of all messages (containing this property) that are associated with each endpoint address.
- **KBPS** — provides an indication of the data transmission rate in kilobytes per second.
- **Duration** — specifies the delta between the **Start** and **End** time values.
- **Start** — specifies the time at which the first message in an endpoint address pair set began.
- **End**— specifies the time at which the last message in an endpoint address pair set ended.
- **K** — a chart constant to facilitate calculation of the KBPS values.
- **BPS** — provides an indication of the data transmission rate in bytes per second.

Using the Top Talkers Layout

To enable you to focus on message volume and payload byte count at different layers, you can set a **Viewpoint** such as **TCP** or **Ethernet**. To apply a **Viewpoint**, click the **Viewpoint** drop-down list on the Filtering toolbar and then select the **TCP** or **Ethernet** item. You will notice that values such as the message count, payload byte count, and data transmission rates may change in the **Layout** as you set these **Viewpoints**.

From this **Layout**, you can obtain statistics such as the following to assist in the troubleshooting process:

- Message volume per endpoint address pair.
- Top bandwidth/data consumers
- Traffic density in terms of payload byte count per endpoint address pair
- Data transmission rates for each endpoint address pair

> **NOTE**
>
> You can double-click any IP **AddressPair** set and display all the messages for a particular set in the **Analysis Grid** viewer for further analysis.

See Also

Top Talkers Top 20

Top Talkers Top 20

2 minutes to read

The built-in **Top Talkers Top 20** view **Layout** for **Charts** provides a summary of the endpoint conversation address pairs that had the top message volume in a set of trace results. The summary data is provided in a **Bar** element visualizer component for the top 20 talkers in a trace and displays the endpoint conversation address pairs and the traffic volume for each pair in a label to the right of the bar elements. A scale is also provided to the far right of the Bar element visualizer that indicates the relative percent volume of messages for each endpoint address pair bar with respect to the overall top 20 talker message count.

Using the Top Talkers Top 20 Chart

To enable you to focus on message volume at different layers, you can set a **Viewpoint** such as **TCP** or **Ethernet**. To apply this **Viewpoint**, click the **Viewpoint** drop-down list on the Filtering toolbar and then select the **TCP** or **Ethernet** item. You will notice that the message count value for the address pair bar elements may change in the **Layout** as you set these **Viewpoints**.

From this **Layout**, you can obtain statistics such as the following to assist in the troubleshooting process:

- Message volume per endpoint address pair
- Top bandwidth/data consumers
- An at-a-glance assessment of how message volume is distributed across the timeline of a trace, from where you can easily identify the top conversations and the endpoints involved.

> **NOTE**
>
> You can double-click any IP address pair set and display all the messages for that set in the **Analysis Grid** viewer for further analysis.

See Also

Top Talkers

Netlogon Category

2 minutes to read

The built-in viewer **Layouts** for **Charts** that are provided with Message Analyzer are accessible from the locations specified in the Chart Viewer Layouts topic. The viewer **Layouts** for **Charts** that are included in the **Netlogon** category of the **Message Analyzer Chart View Layouts** asset collection Library consist of the following, as described in this section:

Netlogon Message Types

Netlogon Message Types

3 minutes to read

The **Netlogon Message Types** view **Layout** for **Charts** enables you to obtain a high-level summary view of specific data from a Netlogon.log file that depicts the relative percentage of message volumes for each message type in the log. The **Layout** uses a **Pie** chart visualizer component where the volume for each message type is represented by a slice of the chart. The legend to the right of the **Pie** chart exposes each message type in a list of nodes for quick correlation with the message volumes represented in the **Pie** chart.

> **NOTE**
>
> A Netlogon.log file will be parsed only if you select the **Netlogon** configuration file in the **Text Log Configuration** drop-down list of the **New Session** dialog for a Data Retrieval Session prior to loading the data into Message Analyzer. Otherwise, no data will display in the **Netlogon Message Types** view **Layout**.

More Information

To learn more about working with text-based .log files, see Opening Text Log Files.

Using the Netlogon Message Types Layout

This **Layout** enables you to assess at-a-glance the distribution of message type volumes in a Netlogon.log file so that you can quickly focus on message types of interest. For example, you will likely want to review any log entries that contain error or diagnostic information first and this would mean looking at the messages of type CRITICAL and DIAGNOSIS. You might also review MAILSLOT messages for records of client and server communications and PERF messages that can include performance counter information related to setting up client-server sessions and the number of authentication timeouts that have occurred.

When your mouse hovers over any slice in the **Pie** chart, a tool tip displays a log message type and a value that represents the percent volume for the message type out of the total volume of all message types in the log. At the same time, the message type is highlighted in the chart's message type legend. Note that this legend is interactive, as is the **Pie** chart itself, and either of these enable you to drive the display of common message types into a new instance of the **Analysis Grid** viewer by double-clicking a legend node or a slice in the **Pie** chart, respectively, for further analysis of a particular log message type.

With these capabilities, you can quickly expose the data of different message types, which is very convenient when you need to detect errors and other important information that is buried in a large log file.

Interactive Analysis

This **Layout** for the **Chart** viewer is intended to work with the **Netlogon Layout** for the **Analysis Grid** viewer and the **Netlogon Group by Message Type Layout** for the **Grouping** viewer to create an integrated and interactive analysis environment. You will be able to correlate the data most effectively if you have these viewers and **Layouts** displayed. Note that these viewers and **Layouts** are configured in the **Netlogon Logs Profile** and will display after you load data from a Netlogon.log file, provided that you enabled this **Profile** on the **Profiles** tab of the **Options** dialog (accessible from the global Message Analyzer **Tools** menu).

As an additional example of interactively driving the display of messages, if you select any **msgtype** group in the **Grouping** viewer, you can display all the messages associated with a particular message type in the **Analysis Grid** viewer for further analysis, which includes reviewing **Summary** information and message **Details**.

More Information

To learn more about interactively analyzing Netlogon log data with the indicated viewing and **Layout** configurations, see the following topics:
Applying and Managing Analysis Grid Viewer Layouts — see the **Netlogon Layout** in this topic.
Grouping Viewer — see the **Netlogon Group by Message Type Layout** in this topic.
Working With Message Analyzer Profiles — see the **Netlogon Logs Profile** in the table of this topic to review a related usage scenario and analysis example and to learn how to manually display the **Grouping** and **Chart** viewers with the **Layouts** defined in this **Profile**.

Networking Category

2 minutes to read

The built-in **Layouts** for the **Chart** viewer that are provided with Message Analyzer are accessible from the locations specified in the Chart Viewer Layouts topic. The **Layout** for the **Chart** viewer that is included in the **Networking** category of the **Message Analyzer Chart View Layouts** asset collection Library is described in the following section:

NTP Time Offset

NTP Time Offset

2 minutes to read

The **NTP Time Offset** view **Layout** for the **Chart** viewer enables you to observe Time Offset data over the timeline of a set of trace results per network conversation, which you can select in a legend to the right of the **Timeline** graphic visualizer component. This will help you understand time offset from the network perspective and to troubleshoot time-related issues. This view **Layout** is accessible from the **Chart** drop-down list that appears when you click **Chart** in the **New Viewer** list on the global Message Analyzer toolbar.

Message Analyzer also provides a **Profile** that displays the **NTP Time Offset** view **Layout** for the **Chart** viewer by default whenever you load a *.cap, *.pcap, *.etl, or *.pcapng file while the **NTP Time Offset Profile** is enabled in the **Options** dialog. This dialog is accessible from the global Message Analyzer **Tools** menu. The **NTP Time Offset Profile** also configures the **NTP Time Offset Layout** for the **Analysis Grid** viewer and the **NTP Source Layout** for the **Grouping** viewer. However, you will need to manually select these **Layouts** as the **Default** in the respective **Layout** lists for these viewers in the **New Viewers** drop-down list on the global Message Analyzer toolbar. Note that these viewers are integrated and interactive so that you can correlate the data through the different viewing perspectives that they create.

> **NOTE**
>
> To obtain meaningful results with the **NTP Time Offset** view **Layout**, the above supported file types must contain time offset data that is captured by the **NTP** module.

More Information

To learn more about the Message Analyzer **Profiles**, see Working With Message Analyzer Profiles.

Common Category

2 minutes to read

The built-in viewer **Layouts** for **Charts** that are provided with Message Analyzer are accessible from the locations specified in the Chart Viewer Layouts topic. The viewer **Layouts** for **Charts** that are included in the **Common** category of the **Message Analyzer Chart View Layouts** asset collection Library consist of the following, as described in this section:

Perfmon Log (.blg)

Perfmon Log (.blg)

2 minutes to read

The **Perfmon Log** view **Layout** for **Charts** enables you to display data from a Performance Monitor log and utilize some of Message Analyzer capabilities to manipulate and analyze the data whenever you load data from a *.blg log file. This **Layout** provides a main display with a graphic representation of performance counter data along with a legend of counters and an adjustable time window for zooming into data points. It displays a related set of messages after you double-click a line of performance counter data for further details.

Interactive Analysis
This **Layout** for the **Chart** viewer is intended to work with the **Perfmon Log Layout** for the **Analysis Grid** viewer and the **Perfmon Log Layout** for the **Grouping** viewer to create an integrated and interactive analysis environment. You will be able to correlate the data most effectively if you have these viewers and **Layouts** displayed. Note that these viewers and **Layouts** are configured in the **Perfmon Logs Profile**. The **Perfmon Log Layouts** for the **Chart** and **Grouping** viewers both display by default after you load data from a *.blg file, provided that you enabled this **Profile** on the **Profiles** tab of the **Options** dialog (accessible from the global Message Analyzer **Tools** menu). However, you will need to manually display the **Analysis Grid** viewer with the **Perfmon Log Layout**.

Consider the following as an example of interactively driving the display of messages into different viewers from the **Grouping** viewer, which contains the following groups to organize the data:

- **Machine**
- **Instance**
- **Counter**

Under any **Instance** group, you can click a **Counter** and display that result in the **Perfmon Log Layout** of the **Chart** viewer. In addition, you can double-click the resulting counter data line and display the logged data in an associated set of messages in a new instance of the **Analysis Grid** viewer. Otherwise, if the **Analysis Grid** viewer is already open in the session, it will simply be updated with the same set of messages.

See Also

Working With Message Analyzer Profiles

File Sharing Category

2 minutes to read

The built-in viewer **Layouts** for **Charts** that are provided with Message Analyzer are accessible from the locations specified in the Chart Viewer Layouts topic. The viewer **Layouts** for **Charts** that are included in the **File Sharing** category of the **Message Analyzer Chart View Layouts** asset collection Library consist of the following, as described in this section:

SMB File Stats
SMB Reads and Writes Bytes Sent
SMB Reads and Writes Bytes/Second
SMB/SMB2 Service Performance
SMB Top Commands
SMB Top Talkers
SysLog Levels

SMB File Stats

3 minutes to read

The **SMB File Stats** view **Layout** for **Charts** provides some basic file access statistics. This **Layout** provides this information in a **Table** grid visualizer component from which you can obtain the following types of data for troubleshooting file access operations:

- The duration of file or folder access operations
- The total bytes for each operation
- The data transmission rate

Using the SMB File Stats Layout

With the statistics provided by this **Layout**, you may be able to determine whether there are bottlenecks occurring during file access operations, network latency issues, or perhaps a poorly responding server. This **Layout** provides the following columns of information:

- **FileName** — a sortable column of data that specifies the name of the file or folder being accessed by the SMB protocol.
- **Duration** — for each row of data, this column provides the time expired during file or folder access operations, in values that resolve to 7 decimal places.
- **Bytes** — for each row of data, this column indicates the total number of bytes for each file or folder access operation.
- **BPS** — for each row of data, this column specifies the data transmission rate in bytes-per-second. The **BPS** values are calculated by dividing the **Bytes** values by the **Duration**.
- **Start** — for each row of data, this column provides a low-resolution timestamp that reflects when the transaction began.
- **End** — for each row of data, this column provides a low-resolution timestamp that reflects when the transaction ended.

> **NOTE**
>
> The **Duration** values are extrapolated from the **Start** and **End** timestamps.

Interactive Analysis

You can also interactively drive the display of messages into the **Analysis Grid** viewer by double-clicking any row of data in the **SMB File Stats** view **Layout**. This results in specific messages in a standard environment for detailed analysis that includes quick access to diagnosis errors and top-level messages that encapsulate message stacks, fragments, and Operations, in addition to the data that is available from the default column **Layout** of the **Analysis Grid** viewer. However, you might also consider using other **Layouts** for the **Analysis Grid** viewer to expose different data sets related to the SMB protocol such as the **File Sharing Perf SMB2/SMB** and **File Sharing SMB/SMB2** view **Layouts**. A particularly useful viewing configuration for analyzing SMB data might be to employ the viewers and **Layouts** that are specified in the table that follows.

Table 13. Interactive Viewing Configuration Example for SMB Data

VIEWER	LAYOUT
Analysis Grid	File Sharing Perf SMB2/SMB
Grouping	File Sharing SMB/SMB2
Chart	SMB File Stats

With this viewing configuration, you can redock the **SMB File Stats** view **Layout** for **Charts** next to the **Analysis Grid** viewer, and then select group nodes in the **Grouping** viewer to simultaneously drive the display of associated messages in the **Analysis Grid** viewer and related rows of data in the **SMB File Stats** view **Layout**. Conversely, you can also select a row of data in the **SMB File Stats** view **Layout** and automatically detect and highlight the messages associated with a particular SMB file operation in the **Analysis Grid** viewer. In addition, message fields and values in the **Details Tool Window** snap to whatever message selections are made through these interactions. In each case, you can obtain focused analysis environments for correlating different aspects of the data.

More Information

To learn more about interactively analyzing SMB data with the indicated viewing and **Layout** configurations, see the following topics:

Applying and Managing Analysis Grid Viewer Layouts — see the **File Sharing Perf SMB2/SMB Layout** in this topic.

Grouping Viewer — see the **File Sharing SMB/SMB2 Layout** in this topic.

Working With Message Analyzer Profiles — see the **File Sharing SMB Perf Profile** in the table of this topic to review a related usage scenario and brief analysis example.

SMB Reads and Writes Bytes Sent

2 minutes to read

The **SMB Reads and Writes Bytes Sent** view **Layout** for **Charts** provides a summary of byte values associated with SMB/SMB2 read, write, and open requests for file operations across a set of trace results. The summary data is presented in a **Table** grid visualizer component that provides an overview of the average and maximum request size in bytes for the SMB/SMB2 message/s associated with each file operation. The table contains the following columns of data:

- **FileName** — specifies the name and path of files upon which an SMB operation was performed.
- **Count** — specifies the number of SMB/SMB2 request messages associated with an operation.
- **Average** — calculates the average size in bytes of the messages sent in an operation.
- **Max** — specifies the highest message size in bytes among all messages sent in an operation.

This **Layout** also has a view **Filter** applied in the background that allows the indicated types of SMB or SMB2 requests to pass while blocking all others, in order to create focus on specific SMB/SMB2 operations.

Using the SMB Reads and Writes Bytes Sent Layout

You can create an interactive analysis context by driving the display of data into the **Analysis Grid** viewer. For example, you can double-click any row of data in the **SMB Reads and Writes Bytes Sent** Table to display the associated messages in a new instance of the **Analysis Grid** viewer for further assessment of **Summary** information and message **Details**.

If you redock the **SMB Reads and Writes Bytes Sent** view **Layout** so that it displays next to the **Analysis Grid** viewer with its default **Layout** in which your original data appears, you can then correlate the data in the **SMB Reads and Writes Bytes Sent** Table with the corresponding messages in the **Analysis Grid** viewer, by simply clicking a row in the Table. Each time you select a Table row, you will have immediate access to **Details Tool Window** data for the corresponding message that is automatically selected and highlighted in the **Analysis Grid**.

You might also consider displaying the **Analysis Grid** viewer with its **File Sharing SMB/SMB2 Layout** to create a grouped viewing configuration that organizes the SMB/SMB2 data by **SessionId** at top-level, **TreeId** at the first nested level, and **FileName** nested below that. You could also drag-and-drop the **FileName** label above the group message nodes in the **Analysis Grid** all the way to the leftmost position in the label order to reorganize the group display with **FileName** nodes at top-level and then use the **Expand All Groups** context command for each group to expose all the messages. You can then easily correlate **FileName** data in the **SMB Reads and Writes Bytes Sent** view **Layout** with corresponding messages in the **Analysis Grid** viewer by simple selection of Table rows.

The advantage of this viewing configuration is that it enables you to correlate the nested **SessionId** and **TreeId** data with the **FileName** row that you select in the **SMB Reads and Writes Bytes Sent** Table. Note that you can also sort the **FileName** column of the **SMB Reads and Writes Bytes Sent** view **Layout** to enable faster location of specific file names.

SMB Reads and Writes Bytes-Second

3 minutes to read

The **SMB Reads and Writes Bytes/Second** view **Layout** for **Charts** enables you to obtain a quick assessment of Server Message Block (SMB) protocol performance from a graphic **Timeline** visualizer component. For example, from the **SMB Reads and Writes Bytes/Second** view **Layout**, you can obtain statistics that reflect the network bandwidth being consumed, in **Bytes per Second**, by the file access/sharing activities of the SMB protocol.

> **NOTE**
>
> The **SMB Reads and Writes Bytes/Second** view **Layout** supports each of the SMB, SMB2, and SMB3 protocol versions.

This **Layout** also has a view **Filter** applied in the background that allows SMB Read, Write, and Open requests to pass while blocking all others, in order to create focus on specific SMB operations.

Using the SMB Reads and Writes Bytes/Second Layout

The main graphic display of this **Layout** provides data in the X-Y coordinate domain, where the X-axis is calibrated in time and the Y-axis is calibrated in bytes/second. You can obtain the following statistics from the **SMB Reads and Writes Bytes/Second** view **Layout** for analysis purposes:

- The name of the file being accessed by the SMB protocol, as indicated in the Legend of the **Layout** and in a tool tip.
- The **Bytes/Second** rate for each SMB request message, as indicated by a Y-axis value and a tool tip.
- The count of SMB request messages sent, as indicated by message nodes that appear when you select a line of data.

Note that you can selectively display or hide data lines associated with specific files for which SMB operations were performed, by respectively selecting and unselecting check box nodes in the file name Legend that appears in the right-hand section of the **Layout**. For each node that you select in the Legend, a line of data appears in the main graphic display at a particular point in the X-axis timeline. For each line of data, there can be one or more data points that appear after you single-click the line, where information for each data point displays in a tool tip when you hover over it with your mouse. This tool tip information consists of the following:

- **File name** — the name of the file upon which an SMB operation was performed.
- **Time** — the time at which the SMB requests occurred.
- **Value** — the bytes/second rate for the SMB request message that is represented as a data point in the data line. This tool tip value is also reflected by a corresponding Y-axis value.
- **Count** — the number of messages associated with a particular data point.

Using the Adjustable Time Windows and Presets

The **SMB Reads and Writes Bytes/Second** view **Layout** has preset **Zoom** values that adjust your view of the data so that you can examine message activity and bytes/second rates within specific windows of time, which includes **1s**, **5s**, **30s**, and **1ms** window presets. After you apply one of these presets, you can return to the original display showing all data lines by clicking the **All** preset. You can also adjust the time-window slider controls in order to zoom into the SMB Read, Write, or Open request activity that occurred in any time slot that you configure.

Interactively Correlating the Data

If you want to display only the SMB requests and supporting message stack that comprised a particular SMB file operation, you can double-click a data line in the visualizer component used in the **SMB Reads and Writes Bytes/Second** view **Layout** and Message Analyzer will display all the associated SMB Operations in a new instance of the **Analysis Grid** viewer. This action provides you with immediate access to **Details Tool Window** data so that you can review message field data and other details for messages that you select in the **Analysis Grid** viewer.

SMB-SMB2 Service Performance

2 minutes to read

The **SMB/SMB2 Service Performance** view **Layout** for **Charts** provides a summary of the top SMB command types in a set of trace results, the number of calls made by each command, along with the minimum, average, and maximum values for the service response time and elapsed time for each command. The summary information is provided as tabular data in the grid visualizer format and contains the following columns of information:

- **Type** — specifies the SMB/SMB2 command type by name, for example `SMB2.VirtualOperations.Create`.
- **Calls** — specifies the number of calls made by each SMB/SMB2 command type.
- **Avg SRT** — averages the **ResponseTime** across all calls made by each different command type, for example `SMB2.VirtualOperations.Read`, `SMB2.VirtualOperations.Write`, and so on.

  The service response time (SRT) is the time differential between an SMB request sent to a server and the first server response to that request, as measured by the **ResponseTime** global annotation that is found in **Field Chooser**. A long SRT can be an indication of a slowly responding server application.
- **Min SRT** — indicates the minimum response time detected across all calls made by each different command type.
- **Max SRT** — indicates the maximum response time detected across all calls made by each different command type
- **Avg Elapsed Time** — averages the elapsed time across all calls made by each different command type, for example, `SMB2.VirtualOperations.Close`, `SMB2.VirtualOperations.QueryInfo`, and so on.

  The elapsed time is the time differential between an SMB request sent to a server and the last response to that request, which signals that all message fragments associated with the operation have been received by the requesting computer. A relatively long elapsed time value for an operation can be an indication of possible network latency issues.
- **Min Elapsed** — indicates the minimum elapsed time detected across all calls made by each different command type.
- **Max Elapsed** — indicates the maximum elapsed time detected across all calls made by each different command type

Using the SMB/SMB2 Service Performance Layout

To maximize the use of the **SMB/SMB2 Service Performance** view **Layout**, the **SMB/SMB2 Viewpoint** is set by default for this **Layout** in the **Viewpoints** drop-down list on the Filtering toolbar. This enables you to view your data from the perspective of the SMB/SMB2 protocol with no layers above it. This also helps you to further isolate the SMB/SMB2 command data along with the associated service response time and elapsed time values, which can provide the previously specified troubleshooting indications.

SMB Top Commands

2 minutes to read

The **SMB Top Commands** view **Layout** for **Charts** enables you to obtain a high-level summary view that depicts the relative distribution of traffic volume, from the highest to the lowest volume, for SMB commands in a set of trace results. The **Layout** uses a **Bar** element visualizer component that provides a label to the left of each bar element to display the type of SMB command. Each horizontal graphic bar element in this **Layout** extends to a certain length corresponding to the command volume level and is tagged with a number that represents the total number of a particular type of SMB command. To the right of each bar element there is also a label that indicates what percent of the total number of SMB commands each bar element represents. In the bottom right sector of the **Layout**, the total number of SMB commands is displayed for the current set of trace results.

Note that this **Layout** has an **SMB/SMB2 Viewpoint** applied by default so that you can isolate SMB and SMB2 messages and create a focused analysis environment with no Application Layer nuances above those messages.

Using the SMB Top Commands Layout

The summary data of this **Layout** enables you to quickly evaluate the SMB commands that are consuming the most bandwidth, as indicated by the longer length bar elements. For example, you could have SMB commands such as **QueryInfo**, **Create**, **Read**, **Close**, **SessionSetup**, **Negotiate**, and others displaying in the **SMB Top Commands** view **Layout** at different volume levels. You might be interested in investigating the source and destination computers that are generating a high volume of Operational messages for a particular type of SMB command. To do this, you can interactively drive the display of messages into a new instance of the **Analysis Grid** viewer by double-clicking any bar element of interest in the **Layout**. Thereafter, you can review the following in the **Analysis Grid** viewer for messages that comprise a particular type of SMB command:

- The IP addresses of the source and destination computers.
- Command status, which includes success and failure indications, along with other information in the **Summary** column.
- Message field names and values in the **Details Tool Window**.

See Also

SMB Top Talkers

SMB Top Talkers

3 minutes to read

The **SMB Top Talkers** view **Layout** for **Charts** provides statistical summary data for the top IP conversations in a set of trace results. It enables you to obtain statistics such as the message count, conversation payload length in bytes, data transmission rate, conversation duration, and others, for the top IP conversations (denoted by address pairs) in a set of trace results. The **Layout** uses a **Table** grid visualizer component that provides this information in the following columns of data:

- **AddressPair** — specifies the IP address of the computers that participated in an IP conversation.
- **Count** — specifies the number of messages that were exchanged in each conversation.
- **Bytes** — based on the **PayloadLength** property of any module that defines this field, the values in this column specify the total payload byte volume of all messages (containing this property) that are associated with each IP conversation.
- **KBPS** — provides an indication of the data transmission rate in kilobytes-per-second for the messages in an IP conversation.
- **Duration** — specifies the delta between the **StartTime** and **EndTime** values.
- **StartTime** — specifies the time at which the first message in a conversation group began.
- **EndTime** — specifies the time at which the last message in a conversation group ended.
- **K** — a chart constant to facilitate calculation of the KBPS values; can be ignored.
- **BPS** — provides an indication of the data transmission rate in bytes-per-second for the messages in an IP conversation.

This **Layout** also contains an `SMB or SMB2` view **Filter** that is automatically applied in the background in order to focus the results on SMB and SMB2 messages only.

Using the SMB Top Talkers Layout

The statistics you can obtain from this **Table** visualizer component that are useful in troubleshooting the SMB/SMB2 protocols include:

- Message volume per conversation
- Top bandwidth consumers
- Data transmission rates
- The time taken to complete conversations

Interactive Analysis

This **Layout** is intended to work with the **SMB Flat** view **Layout** of the **Analysis Grid** viewer and the **File Sharing SMB/SMB2** view **Layout** of the **Grouping** viewer to create an integrated and interactive analysis environment. You will be able to correlate the data most effectively if you have these viewers and **Layouts** displayed. Note that these viewers and **Layouts** are configured in all four of the **File Sharing SMB Profiles** and will display after you load data from a .cap, .etl, .pcapng, or .pcap file, provided that you enabled the appropriate **Profile** on the **Profiles** tab of the **Options** dialog (accessible from the global Message Analyzer **Tools** menu).

For example, to interactively analyze SMB and SMB2 messages, you can drive the display of these messages into the **Analysis Grid** viewer by double-clicking any row of data for a particular conversation in the **SMB Top Talkers** view **Layout**. Thereafter, you can review SMB status information for any message in the **Summary** column of the **Analysis Grid** viewer along with message fields and values in the **Details Tool Window**. The **File Sharing SMB/SMB2 Layout** for the **Grouping** viewer also enables a greater interactive context with additional enhancements to the analysis process.

More Information

To learn more about interactively analyzing SMB/SMB2 data with the indicated viewing and **Layout** configurations, see the following topics:

Applying and Managing Analysis Grid Viewer Layouts — see the **SMB Flat Layout** in this topic.

Grouping Viewer — see the **File Sharing SMB/SMB2 Layout** in this topic.

Working With Message Analyzer Profiles — see the **File Sharing SMB Profile** in the table of this topic to review a related usage scenario and analysis example and to learn how to manually display the **Grouping** and **Chart** viewers with the **Layouts** defined in this **Profile**.

See Also

SMB Top Commands

SysLog Levels

3 minutes to read

The **SysLogLevels** view **Layout** for **Charts** enables you to quickly assess the relative distribution of the log entry volumes per information **level** value that exist in a SambaSysLog file. The **Layout** uses a **Bar** element visualizer component that provides a label to the left of the bar elements to display the value of associated information **levels** in a SambaSysLog. Each horizontal graphic bar element in this **Layout** extends to a certain length that corresponds to the log entry volume for a specific information **level**. Each bar element is also tagged with a number that appears to the right of the element to represent the total number of entries that contain a particular **level**. To the right of the bar element display, there is set of labels that indicate what percent of the total number of log entry **levels** each bar element represents. Also, in the bottom right sector of the **Layout**, the total number of log entries with level designations is displayed.

> **NOTE**
> A SambaSysLog.log file will be parsed only if you select the **SambaSysLog** configuration file in the **Text Log Configuration** drop-down list of the **New Session** dialog for a Data Retrieval Session prior to loading the data into Message Analyzer. Otherwise, no data will display in the **SysLogLevels** view **Layout**.

More Information
To learn more about working with text-based .log files, see Opening Text Log Files.

Using the SysLogLevels Layout

With the **SysLogLevels** view **Layout** for **Charts**, you can instantly assess the areas in your log that had the most critical **levels**, which can indicate an initial direction in which further investigation might proceed. SambaSysLog **levels** typically consist of the following. As expected, errors and warnings are usually the most critical to review, as described in "Interactive Analysis" ahead:

- 0 — Error
- 1 — Warning
- 2 — Notice
- 3 — Information
- 4+ — Debug

Interactive Analysis
The **SysLogLevels** view **Layout** is intended to work with the **SysLog** view **Layout** of the **Analysis Grid** viewer and the **SysLog** view **Layout** of the **Grouping** viewer to create an integrated and interactive analysis environment. You will be able to correlate the data most effectively if you have these viewers and **Layouts** displayed. Note that these viewers and **Layouts** are configured in the **Samba Logs Profile** and will display after you load data from a SambaSysLog.log file, provided that you enabled this **Profile** on the **Profiles** tab of the **Options** dialog (accessible from the global Message Analyzer **Tools** menu).

For example, to interactively analyze SambaSysLog data, you can drive the display of log entries that contain **level** designations into the **Analysis Grid** viewer by double-clicking any bar element for a particular **level** in the **SysLogLevels** view **Layout**. Thereafter, you can review status information and other content for any log entry in the **Summary** column of the **Analysis Grid** viewer along with log entry fields and values in the **Details Tool**

Window. If you want to isolate log entries that contain an error or warning, you might apply a view **Filter** to the **Analysis Grid** viewer such as `*Summary contains "level=0"` or `*Summary contains "level=1"`, respectively. You can also quickly isolate **level** data from high volume SambaSysLog.log files with the **Grouping** viewer, as described immediately below.

The **SysLog** view **Layout** for the **Grouping** viewer also enables a greater interactive context with additional enhancements to the analysis process. For example, this **Layout** is organized with **level** values as the top groups with nested **function** groups, and nested **source_file** groups beneath that. The advantage of this viewing configuration is that you can isolate the log entry data to the top group (the Samba information **level**), then to the first nested group (the Samba **function** that wrote the log entry), and finally to the second nested group (the Samba **source_file** that contains the function that logged an entry). As you select these groups, the log entries associated with them display in the **Analysis Grid** viewer for a focused assessment of details. This grouped configuration enables you to streamline and prioritize your investigation based on the **level** values, which is a good starting point from where you can determine, in a hierarchical manner, the functions and source code that are associated with the most critical **levels**.

More Information

To learn more about interactively analyzing SambaSysLog data with the indicated viewing and **Layout** configurations, see the following topics:

Applying and Managing Analysis Grid Viewer Layouts — see the **SysLog Layout** in this topic.

Grouping Viewer — see the **SysLog Layout** in this topic.

Working With Message Analyzer Profiles — see the **Samba Logs Profile** in the table of this topic to review a related usage scenario and analysis example and to learn how to manually display the **Grouping** and **Chart** viewers with the **Layouts** defined in this **Profile**.

Interaction Viewer

4 minutes to read

Message Analyzer provides the **Interaction** diagram, which is currently a preview feature. If you wish to use the **Interaction** viewer, you will need to select its check box on the **Features** tab of the **Options** dialog, which is accessible from the global **Tools** menu, and then restart Message Analyzer. It will then be available for selection in the **New Viewer** drop-down list on the global Message Analyzer toolbar.

Understanding the Interaction Diagram

The **Interaction** diagram is an alternate data visualizer that you can select when analyzing your message data. The **Interaction** diagram provides several viewer formats with different graphic presentations of the computer nodes and endpoints that exchanged messages within the time boundaries of a trace. It also specifies the IP address of the corresponding source and destination nodes or endpoints. The IP communications between source and destination nodes are depicted in the following viewer formats:

- **Swimlane** — the default viewer format that displays when you launch the **Interaction** viewer. Displays the IP communications as rows or lines that interconnect source and destination nodes. **MessageNumber** and **Timestamp** data is included for identifying messages and the time of capture.
- **Mapping** — provides a mapped graphic presentation that displays color-coded boxes at the intersection of source IP addresses on the X-axis with destination IP addresses on the Y-axis.
- **Diagram** — shows source and destination interactions as lines that are connected between blue colored nodes that vary in thickness, depending on the number of messages they represent. Diagnosis error or informational icons display over connecting lines where errors occurred. IP addresses, protocol types, message count, and/or diagnosis types display when hovering over connection lines or nodes.
- **Chord** — provides a circular display of IP addresses with interconnecting lines between source and destination nodes in the center of the display. Hovering over an interconnecting line causes an informational tooltip to display while all other lines disappear. The tooltip provides the number of selected items, source and destination IP addresses, and the number of messages associated with each interaction.

The figure that follows shows the **Interaction** viewer in the **Swimlane** configuration.

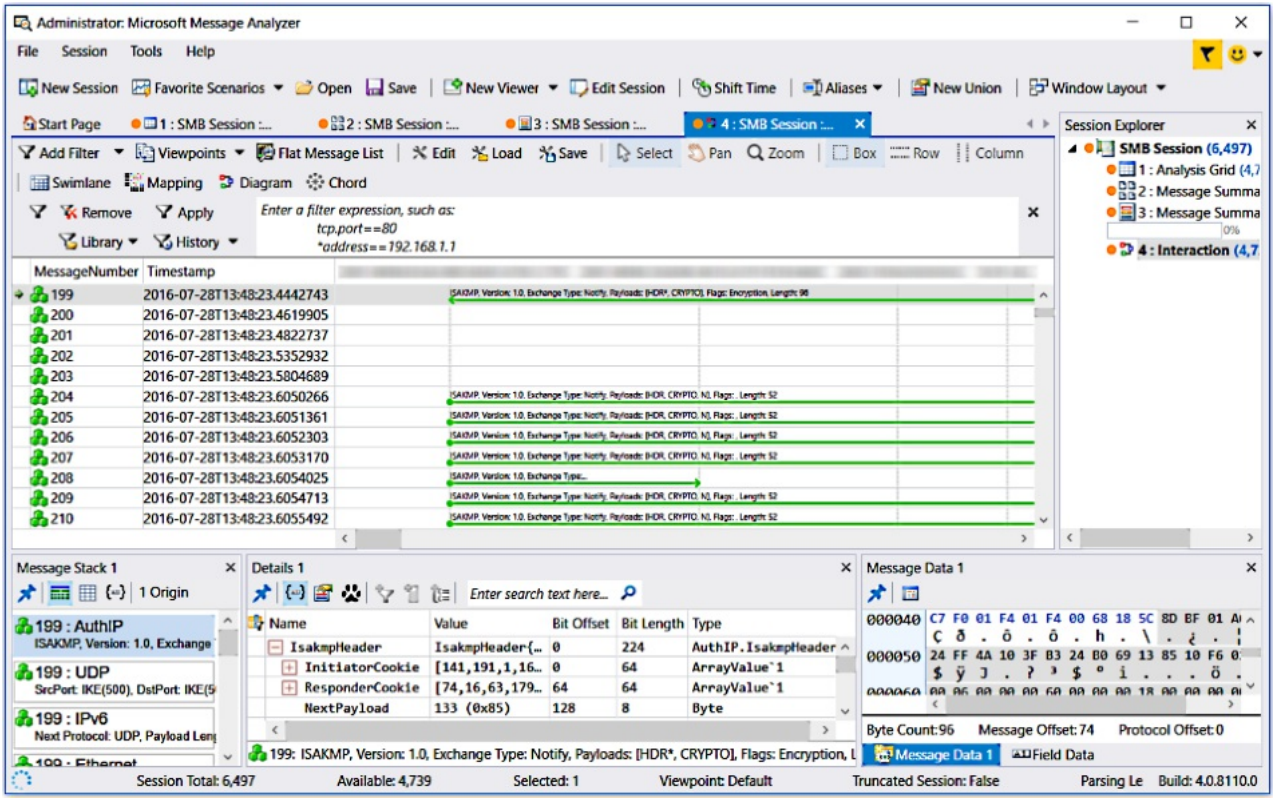

Figure 47: Interaction viewer with Swimlane diagram

You can use the features of the **Interaction** diagram to highlight the communication paths between nodes within the time boundaries of a trace and to identify the relative time within trace boundaries in which network conversations took place. For example, you might use the **Interaction** diagram to expose a map of communications between IP address nodes to identify a busy server, by observing which node is exchanging the most network traffic. You can also observe the direction of traffic flow between nodes.

- **Select** mode — provides the **Box**, **Row**, and **Column** selection formats for viewing data.
- **Zoom** — enables you to zoom into areas in the data display of specific interest.

Using the Toolbar Commands

The **Interaction** diagram provides a toolbar with the following commands, which enable you to perform the indicated tasks:

- **Edit** — click this button to open the **Swimlane Configuration Editor**. Enables you to modify the data display by using the **Field Chooser**. For example, you could add **ModuleName** and/or **ProcessId** fields from the **Global Properties** node of **Field Chooser** for **Source** computers to include this information in the column labels of the **Swimlane** diagram, to enhance your analysis process.
- **Load** — enables you to load data into the current **Interaction** viewer display from a previously saved **Interaction** diagram file in *.maidconfig format.
- **Save** — enables you to save data from the current **Interaction** viewer display to a *.maidconfig file.
- **Select** mode — not configurable. Defaults to the **Box** style selection mode.
- **Pan** — not used.
- **Zoom** mode — not used.

The **Mapping**, **Diagram**, and **Chord** viewer formats provide the following additional toolbar items, which are further described in the Gantt Viewer topic:

Using the Context Menu Commands

The **Interaction** diagram provides two different right-click context menus that display, depending on the **Interaction** viewer format that you select. The context menus appear when you right-click anywhere on a diagram surface. The context menu contains the following commands that perform the indicated operations:

- **Swimlane** viewer format — the context menu commands in this viewer format consist of the following:
  - **Increase Font** — enables you to incrementally increase the size of the diagram display font with each successive selection of this command.
  - **Decrease Font** — enables you to incrementally decrease the size of the diagram display font with each successive selection of this command.
  - **Show Multiline Headers** — enables the display of all lines of text for multiline headers.
  - **Open Selected Messages** — enables you to display one or more selected messages from the **Interaction** diagram in the **Analysis Grid** for further analysis, for instance, to add a view **Filter** and/or group the messages.
- **Mapping**, **Diagram**, and **Chord** viewer formats — the context menu commands in these viewer formats consist of the following:
  - **Mouse Mode** — contains a submenu with the **Select** and **Zoom** options to enable you to change the actions that occur when you click your mouse in the analysis surface. One of these options will always be selected. The **Select** mode facilitates the selection and processing of messages, while the **Zoom** mode initiates zooming action when you click in the analysis surface.
  - **Selection Mode** — provides a submenu that has the following options that you can use as indicated:
    - **Box** — sets the selection mode to the box shape for selecting one or more messages nodes.
    - **Row** — sets the selection mode to enable you to select messages in row orientation.
    - **Column** — sets the section mode to enable you to select messages in column orientation.
  - **Zoom** — contains a submenu that provides several zooming presets that enable you to alter the data presentation to focus on one or more messages.
  - **Settings** — contains a submenu that enables you to change the data presentation format by alternately enabling or disabling various chart elements that include **Fixed Grid**, **Auto Scroll**, and **Display Limit** settings.

Message Summary Tiles Viewer

7 minutes to read

Message Analyzer provides the **Message Summary Tiles** viewer, which is currently a preview feature. If you wish to use the **Message Summary Tiles** viewer, you will need to select its check box on the **Features** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu, and then restart Message Analyzer. It will then be available for selection in the **New Viewer** drop-down list on the global Message Analyzer toolbar.

Message Summary Tiles Viewer Overview

The **Message Summary Tiles** data viewer summarizes important data for any set of trace results. It provides a high-level overview of major trace statistics and important values that you can examine at-a-glance to obtain a quick top-level analysis of results. The default configuration of this viewer contains six tiles of summary information, however, you can configure your own summary tiles by selecting from the preset tile types. Some of the preset tile types require custom configuration; for example, the **Field** type, for which you specify fields, properties, annotations, or other values that you select from the **Field Chooser Tool Window**. Other preset tile types that you must custom configure are the **Filter** and **Sequence** types, which require you to select a predefined asset type or create your own to add them to the tile configuration. The default configuration of the **Message Summary Tiles** viewer is shown in the figure that follows.

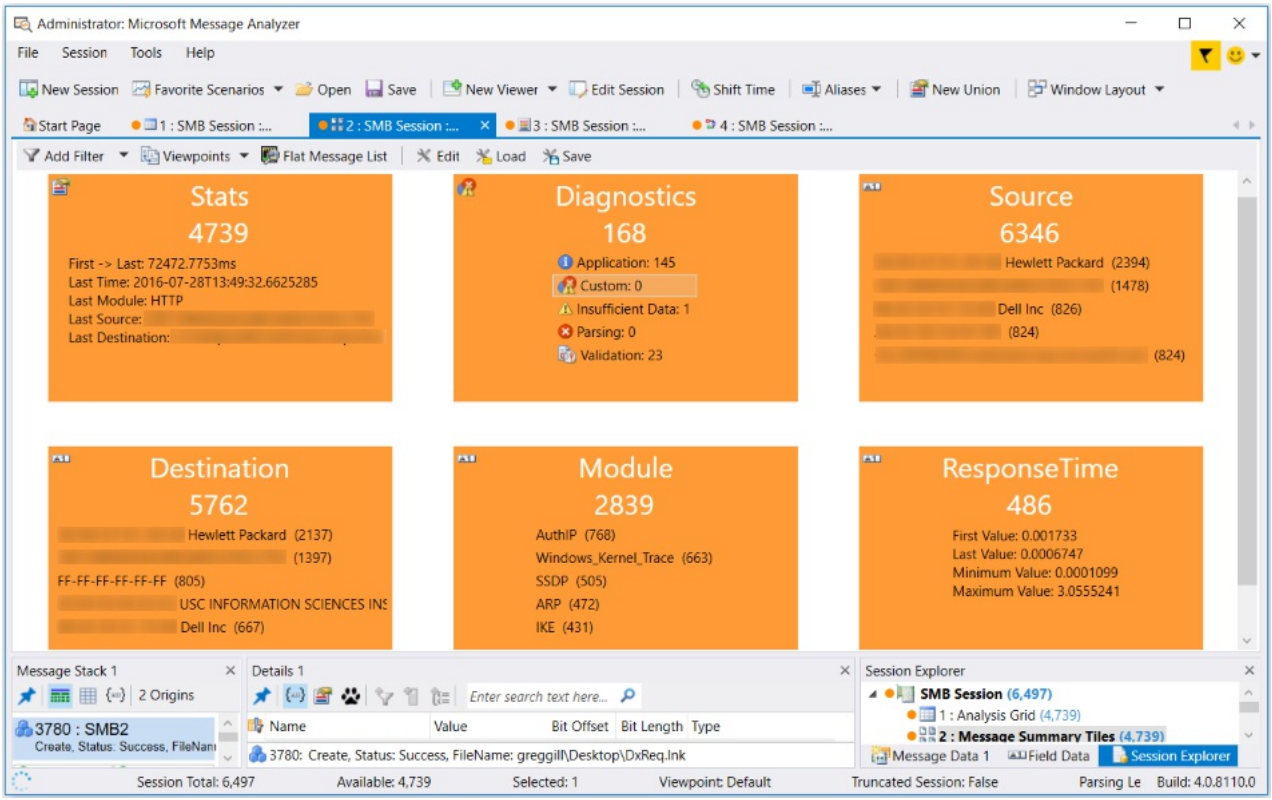

Figure 48: Message Summary Tiles viewer

Isolating Messages for Analysis

Each of the default tiles contain lines of data that are labeled according to predefined fields and you can drill down into the messages that are associated with the details that display on several tiles. You can do this by double-clicking any line of data to display the associated messages in a separate **Analysis Grid** viewer instance that contains only those messages, for further inspection. For example, you might double-click the **Validation** label in the **Diagnostics** tile to view messages for which validation errors occurred, as these can indicate that the

messages of a particular protocol had field values that were out of tolerance or did not meet other specifications when evaluated during the Message Analyzer Runtime parsing process. The only exceptions to double-clicking tile labels to drill down is the **Stats** tile, which always displays all messages whenever you click anywhere in this tile, and the **ResponseTime** tile, which does not currently respond to double-clicking.

Understanding the Default Summary Tiles

The functions of the default summary tiles are described as follows:

- **Stats** — provides an indication below the tile name, of the overall number of messages that were evaluated according to predefined criteria; below that is an overview of the following general data for the current set of trace results:
  - **First->Last** — indicates the duration in milliseconds (ms) of the current trace.
  - **Last Time** — provides the timestamp of the last message in the current trace.
  - **Last Module** — provides the name of the last module in the last message of the current trace.
  - **Last Source** — specifies the source IP address associated with the last message in the current trace.
  - **Last Destination** — specifies the destination IP address associated with the last message in the current trace.

  > **NOTE**
  > Double-clicking anywhere on this tile displays all trace messages in a new **Analysis Grid** viewer instance.

- **Diagnostics** — provides an indication below the tile name, of the overall number of messages that were evaluated according to predefined criteria; below that is an overview of the following diagnostic message data for the current set of trace results:
  - **Application** — specifies the number of messages in the current trace that contain **Application** type diagnostic errors.
  - **Custom** — an alternate type of diagnostic message. Currently a placeholder.
  - **Insufficient Data** — specifies the number of messages in the current trace that contain **Insufficient Data** type diagnostic errors.
  - **Parsing** — specifies the number of messages in the current trace that contain **Parsing** type diagnostic errors.
  - **Validation** — specifies the number of messages in the current trace that contain **Validation** type diagnostic errors.

More Information
To learn more about the meaning of diagnostic messages, see the Diagnosis Category topic.

- **Source** — provides an indication below the tile name, of the total number of values that were evaluated in the top 5 according to predefined criteria; below that is the message count of the specific sender IP addresses with the top 5 highest message volumes for the current set of trace results.
- **Destination** — provides an indication below the tile name, of the overall number of values that were evaluated in the top 5 according to predefined criteria; below that is the message count of the specific recipient IP addresses with the top 5 highest message volumes for the current set of trace results.
- **Module** — provides an indication below the tile name, of the overall number of values that were evaluated

in the top 5 according to predefined criteria; below that is the message count for the specific protocols or modules with the top 5 highest message volumes for the current set of trace results.

- **ResponseTime** — provides an indication below the tile name, of the overall value count for the evaluations performed according to predefined criteria. Below that are the following **ResponseTime** statistics for operations in the current set of trace results that specify the difference in milliseconds between the time a client request message is sent to a server and the time at which the first server response message is received by the client:
  - **First Value** — specifies the **ResponseTime** for the *first* operation that was detected in the current set of trace results.
  - **Last Value** — specifies the **ResponseTime** for the *last* operation that was detected in the current set of trace results.
  - **Minimum Value** — specifies the minimum **ResponseTime** value that was detected in the current set of trace results.
  - **Maximum Value** — specifies the maximum **ResponseTime** value detected in the current set of trace results.

Note These statistics can quickly alert you when there is a problem with poorly performing servers, in terms of responses to client requests.

> **TIP**
> For additional summary information about any of the default summary tiles, hover over the number just below the tile name with your mouse to display a popup window that specifies related evaluation data.

Modifying the Message Summary Tiles Display

You have the option to modify any of the default summary tiles or you can create new tiles by using the **Tile Config** editor to create your own data summary configurations. You can also **Save** your modifications as a Message Summary Viewer Configuration (*.msvcfg) file, that you can **Load** back into Message Analyzer any time thereafter for analysis purposes. You can access the **Tile Config** editor (and the **Load** and **Save** commands) in the Message Analyzer global **Session** menu whenever the **Message Summary Tiles** viewer is displayed. The editor enables you to specify settings such as the level to which **Top Values** will be evaluated; the fields, annotations, or properties you want to use for a data summary through access to the **Field Chooser** window; the tile position in the display, tile **Foreground** and **Background** colors, a textual **Description**, and so on.

Preset Tile Types
The **Tile Config** editor provides a number of preset tile types, which include the following:

- **Stats** — select this preset type if you want to display predefined general statistics in your summary tile. Note that this tile is included in the default configuration of the **Message Summary Tiles** viewer.
- **Diagnostics** — select this preset type if you want to display predefined diagnostics information in your summary tile. This tile is included in the default configuration of the **Message Summary Tiles** viewer.
- **Source** — select this preset type if you want to add predefined **Source** computer information to your summary tile. Note that this tile is included in the default configuration of the **Message Summary Tiles** viewer.
- **Destination** — select this preset type if you want to add predefined **Destination** computer information to your summary tile. This tile is included in the default configuration of the **Message Summary Tiles** viewer.
- **Module** — select this preset type if you want to add other **Module** information to your summary tile. Note

that this tile is included in the default configuration of the **Message Summary Tiles** viewer.

- **ResponseTime** — select this preset type if you want to add predefined **ResponseTime** information to your summary tile. This tile is included in the default configuration of the **Message Summary Tiles** viewer.

> **NOTE**
>
> If you want to add other information based on a specific message field, click the **+** button to display a **(Field)** placeholder item in the **Field Tiles** list, the name for which will change when you select a particular field that you want from the **Field Chooser Tool Window**. You can open **Field Chooser** by first selecting the **(Field)** item and then clicking the **Click to Set** button.

See Also

Field Chooser Tool Window

Message Summary Lists Viewer

2 minutes to read

Message Analyzer provides the **Message Summary Lists** viewer, which is currently a preview feature. If you wish to use the **Message Summary Lists** viewer, you will need to select its check box on the **Features** tab of the **Options** dialog, which is accessible from the global **Tools** menu, and then restart Message Analyzer. It will then be available for selection in the **New Viewer** drop-down list on the global Message Analyzer toolbar.

Understanding the Message Summary Lists Viewer

Message Analyzer enables you to use the **Message Summary Lists** viewer to obtain statistics that reflect key data points across the time line of a set of trace results. These data points provide basic summary information about a trace at-a-glance, as shown in the figure that follows:

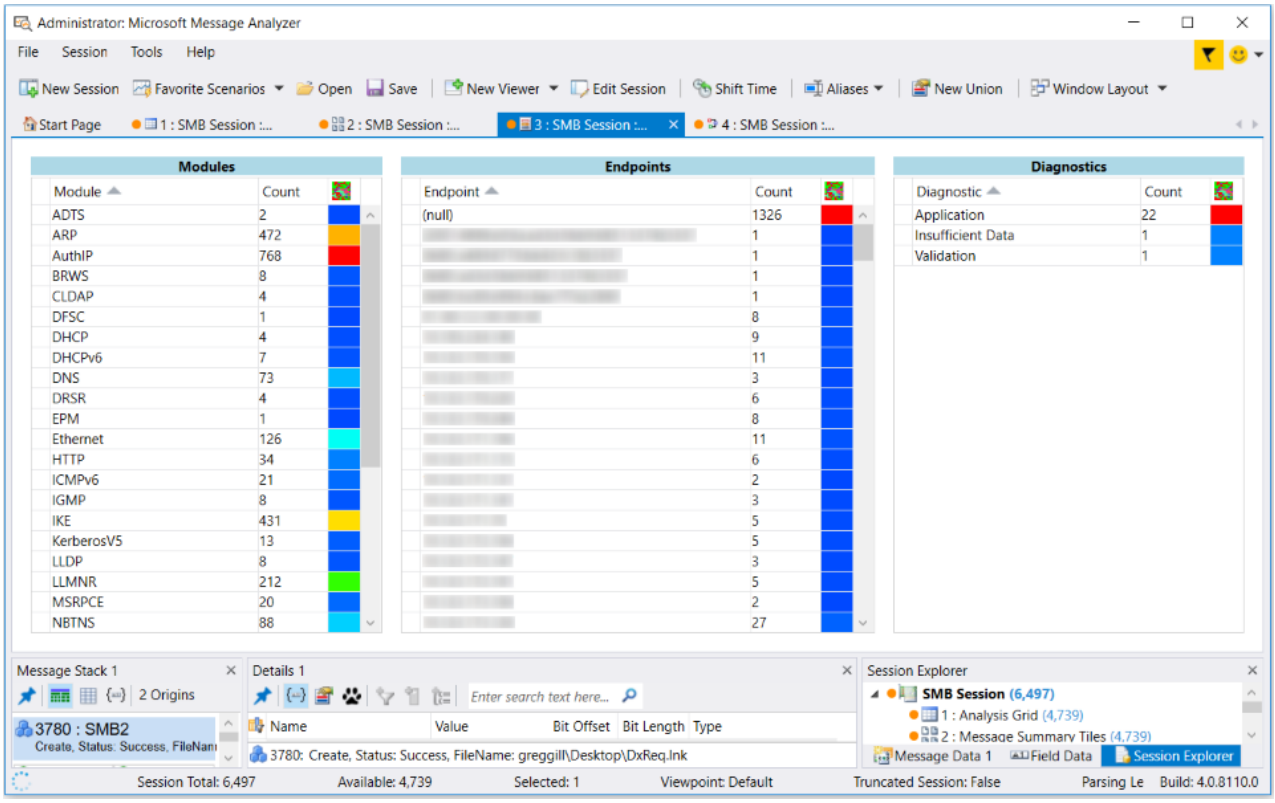

Figure 49: Message Summary Lists viewer

Viewing Statistical Summary Data

The statistics data that you can view is organized into three different categories and within each there are three sortable columns of data. The categories for the statistics data that you can observe with this viewer consist of the following:

> **NOTE**
>
> This viewer does not interactively drive the display of data in any other data viewer or **Tool Window**.

- **Modules** — displays each **Module** that supported a message conversation, along with the total message **Count** for each **Module**, across the time line of a trace.
- **Endpoints** — the address of each **Endpoint** associated with a message conversation, along with the total

message **Count** for each **Endpoint** conversation, across the time line of a trace.

- **Diagnostics** — the total **Count** of **DiagnosisType** messages detected in a set of trace results, which includes **Application**, **Validation**, **Insufficient Data**, and **Parsing** message types.

More Information
To learn more about **Diagnosis** message types, see the Diagnosis Category topic.

You can sort each column in any of the indicated categories to organize the data display differently, which includes sorting the heatmap column. The heatmap column in each category provides a color-coded visual representation of the message volume in each category, as indicated by the **Count** column. Higher volumes correspond with the red shades and lower volumes correspond with the blue shades.

> **TIP**
> To isolate the data that is associated with any selected row of data in the **Modules**, **Endpoints**, or **Diagnostics** categories of the **Message Summary Lists** viewer, right-click a data row and select the **Open Selected Items** command in the context menu that appears. Message Analyzer will then open a new **Analysis Grid** viewer session tab to display the data associated with the row you selected, for further examination.

See Also

Message Summary Tiles Viewer

Perfmon Viewer

6 minutes to read

Message Analyzer provides the **Perfmon Viewer**, which is currently a preview feature. If you wish to use the **Perfmon Viewer**, you will need to select its check box on the **Features** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu, and then restart Message Analyzer. It will then be available for selection in the **New Viewer** drop-down list on the global Message Analyzer toolbar.

Understanding the Perfmon Viewer

Message Analyzer enables you to view data from **Microsoft Performance Monitor** log files in the *.blg format with use of the **Perfmon Viewer**. You can view this data similarly to the way it appears in **Performance Monitor**. To assess this type of data, you might begin by creating a custom data collector set with **Performance Monitor** that contains one or more performance counters. **Performance Monitor** enables you to select from a set of templates that each specify some performance counter presets, or you can manually create your own counters. For example, you could create a custom data collector set that configures one or more performance counters that monitor and return statistics for various entities such as HTTP services, Event Tracing, Hyper-V adapters, IPv4 performance, SMB operations, TCP performance diagnostics, and so on. After you configure and name a data collector set, it will appear under the **User Defined** node in the **Data Collector Sets** container of the **Performance Monitor** MMC.

After collecting data, you will need to right-click the data collector set name under the **User Defined** node and select the **Stop** command in the context menu that appears, before you can load the output log data into Message Analyzer. You can then view the collector set log file data by loading it into Message Analyzer through a Data Retrieval Session. After you load and display the data, typically in the **Analysis Grid** viewer, you can open the **Perfmon Viewer** from the **Common** category of the **New Viewer** drop-down list on the global Message Analyzer toolbar. To graphically display the performance counter data stored in the log, you will need place a check mark in each **Show** check box of the **Perfmon Viewer** for specific counters that collected data you want to view. The **Perfmon Viewer** specifies a different color code for the data lines of each preset performance counter and also provides controls that enable you to **Zoom** into preset windows of time for analysis.

> **TIP**
> You might consider correlating the performance counter data that you collect with that of another related data source, such as an ETL file, live trace, or Event Log, for an enhanced analysis perspective.

More Information
To learn more about using multiple data sources in a Data Retrieval Session, see Configuring Session Scenarios with Selected Data Sources.

Using the Perfmon Viewer Controls for Analysis

Several controls are available in the **Perfmon Viewer** to facilitate analysis of collected data. To analyze counter data with the **Perfmon Viewer**, you can use the controls and indicators that exist in the lower and upper sectors of the viewer interface, which are referred to as the Counter Table and the Data Grid, respectively. The controls can change the way you display the data to facilitate analysis. For example, you can select which counters will display data, set the ranges of the X and Y axis in the data grid to facilitate focused analysis, and select individual data points in the data grid to analyze the counter statistics associated with a selected data point. The controls and

indicators in each sector of the viewer interface consist of the following:

- **Counter Table** — this table is located in the lower sector of the **Perfmon Viewer**. It consists of a list of the performance counters that recorded data in a *.blg file. The table contains the following columns, some with control features and others with data that is associated with each counter that collected data:
  - **Show** — enables you to select one or more counters so that you can isolate specific data of interest for display in the data grid.
  - **Color** — provides an indicator of the default color for each counter. Also enables you to change the default color for any particular counter by clicking the down arrow in the **Color** column of a specific counter to display the **Color** dialog, from where you can specify a color of choice.
  - **Scale** — specifies the current value of the Y-axis scaling for a particular line of counter data. Click the down arrow in the **Scale** column of a particular counter to display a drop-down list containing preset **Scale** values from which you can select. With this control, you can manipulate the range of the Y-axis, which plots the relative deviations of counter values over time.
  - **Counter** — each row of the counter table displays the name of a counter in this column.
  - **Instance** — provides an indication of the number of instances associated with a counter, such as you might have with processes or threads running with the counter and collecting data.
  - **Object** — specifies the name of the object from which a counter is collecting data.
  - **Computer** — specifies the name of the computer on which a counter is collecting data.
- **Data Grid** — the Data Grid is located in the lower sector of the **Perfmon Viewer**. It displays the data for any selected counter in the X-Y coordinate domain. The X-axis is calibrated in time while the Y-axis is calibrated for the relative deviation of counter values. To focus in on your counter data, you can change the X-axis calibration by setting **Scale** values, as described earlier, and you can change the Y-axis calibration with use of the configuration toolbar described in the list below.

  You can also select data points in the Data Grid by clicking on a particular point in a counter data line. At that time, you will see the associated counter data appear in the **Message Stack Tool Window** of Message Analyzer and the values of associated data fields will display in the **Details Tool Window**. For example, you might review the **Value** field for a selected data point in the **Values** column of the **Details** window.

  > **TIP**
  > Through data point selection in the Data Grid, you may be able to correlate information across multiple related logs that you retrieve with Message Analyzer.

  The Data Grid also contains a context menu with the following commands that appear when you right-click anywhere in the Data Grid surface:

  - **Highlight Selected** — when you select this command, it causes the line of counter data in the grid to be bolded as you select it in the **Show** column of the counter table.
  - **Show Grid Lines** — toggle this command to alternatively show and hide grid lines in the Data Grid.
  - **Show Search/Configuration Bar** — when you select this command, it causes the **Search/Configuration Bar** to display. The search control on this toolbar enables you to specify a search string that reflects a particular value for which you are looking in any column. The configuration portion of this bar enables you to recalibrate the X-axis time range (and format) to **Zoom** in on data for focused analysis. Other controls enable you to select or unselect all counters simultaneously.

Charts (Deprecated)

2 minutes to read

The **Charts (Deprecated)** viewers consist of those **Charts** that were available prior to the release of Message Analyzer version 1.4. You can still use these whenever you want, as they are still fully functional. Many of these **Charts**, or their components, are now represented as view **Layouts** that are accessible from the **Charts** drop-down in the **New Viewer** drop-down list. The paradigm changed to simplify things, so that there is now only one **Chart** viewer and many different types of **Layouts** that each use a single visualizer component such as a bar element or grid. For now, only the most popular of the former **Chart** viewers is included in this section for convenience, as indicated immediately below. It contains a combination of all the visualizer components available rather than a single component, as previously indicated.

Protocol Dashboard

Protocol Dashboard

4 minutes to read

The **Protocol Dashboard** is a **Chart**-style data viewer that enables you to obtain a quick assessment of trace performance from graphic visualizer components that provide top-level summary information. It exists in the **Dashboards** category of your local **Charts** asset collection Library, which is accessible from the locations described in Chart Viewer Layouts. A description of the visualizer components that are contained in the **Protocol Dashboard** follows.

Using the Protocol Dashboard

The **Protocol Dashboard** contains the following visualizer components:

- **Top Level Protocol Summary** grid — consists of a table that provides a summary of the top-level messages retrieved by various modules in a trace, along with the number of top-level messages received from each module.

> **NOTE**
>
> If you double-click a table row, all messages in which the module is either a top-level message or part of message origins will display in a separate **Analysis Grid** viewer tab. Because *all* messages are displayed in a separate viewer tab as indicated, the message count in that viewer tab will differ from the top-level message count specified in the **Top Level Protocol Summary** table. To see only the top-level messages in the new **Analysis Grid** viewer tab, you can apply a view **Filter** that isolates top-level messages by using the backslash symbol ("\") described in Browsing Message Origins, as follows:
>
> `"\<modulename>"`
>
> — where <modulename> is a placeholder for the module or protocol of interest.

- **Top Level Protocol Summary** bar **Chart** — this visualizer component illustrates message data in a linear graphic format that provides an at-a-glance view of the relative volume of the top-level messages received from different source modules over the trace timeline. If you double-click any data bar representing a message source in this **Chart**, the details for the top-level messages represented by that particular data bar are rendered in a separate **Analysis Grid** viewer tab for further analysis.

> **NOTE**
>
> If you double-click a bar in the bar **Chart** visualizer component, you will see the same message count disparity as described in the previous note. Similar to the table grid, you can see the top-level messages by applying a view **Filter** with the backslash symbol ("\"), as previously indicated.

- **Top Level Protocol Summary** pie **Chart** — similar to the **Top Level Protocol Summary** bar **Chart**, this visualizer component also illustrates the relative volume of top-level messages received from different source modules; however, the data is presented in a pie chart configuration instead. In this **Chart**, message volume for source modules is delineated as a percentage of the whole so you can quickly obtain the top protocol usage-percentage in the trace. You can also double-click any source module in the **Top Level Protocol Summary** pie **Chart** to present the associated top-level message details in the Message Analyzer **Analysis Grid** viewer.

NOTE

Both the bar and pie **Chart** visualizer components can provide an at-a-glance indication of the top bandwidth consumers in a trace.

- **Top Level Protocols Over Time** timeline **Chart** — this visualizer component provides a graphic, X-Y axis formatted chart. This component depicts message count as X-axis module lines that indicate the number of messages that were captured for a particular module across the entire trace timeline, with individual message nodes indicating the points in time when messages were captured. You can view the message count at any node for any module by hovering over one with your mouse and also in the legend to the right of the graphic viewer surface. In addition, the Y-axis is calibrated to provide an indication of relative message count for any node.

  In addition, if you double-click a node or a module line, all the top-level messages and origins associated with that module will display in a separate **Analysis Grid** viewer tab for further analysis. You can also focus on the messages from a particular module by selecting them in a legend to the right of the timeline, and clearing selections of the message types you do not want to display. Lastly, this visualizer component enables you to adjust selectable time-window slider controls so that you can do the following:

  - Visually assess the times at which protocol communications occurred within the time range of a trace.
  - Use manual adjustments to drill down into specific time slots for a more granular view of the message activity that transpired there.
  - Use **Zoom** presets to automatically create time windows that enable you to drill down into message activity in various time slots, starting from the beginning of a trace.

    These preset values create time windows that are **100 ms**, **1s**, or **5s** in length. After you specify one of these presets, you can return to the full trace time boundaries by clicking the **All** preset, or you can manually expand the trace boundaries with the time window slider controls.

TIP

If you apply a view **Filter** or a **Time Filter** to the **Protocol Dashboard** viewer in a particular session, messages will be appropriately filtered in this viewer only, while the message data in any other viewer opened in the same session remains unaffected.

Session Data Viewer Options

5 minutes to read

After you use the **New Session** dialog to configure either a Data Retrieval Session that points to the input files that contain the data you want to work with, or a Live Trace Session that you are ready to run, you can choose one of the data viewers previously described in the Data Viewers section to display your data. You can also simply allow Message Analyzer to present your session results in the default **Analysis Grid** viewer, unless you have already set another type of data viewer as the default in the **Options** dialog, as described in Setting the Default Session Viewer. If you want to choose a data viewer for a session, select one from the common **Start With** drop-down list in the **New Session** dialog prior to starting the session. After you click the **Start** button in the **New Session** dialog, Message Analyzer immediately begins to retrieve or capture data and display it in the chosen viewer.

Using the Data Viewing Options

The data viewer that appears in the text box portion of the **Start With** combo box when you open the **New Session** dialog, is the default viewer that is set in the global **Options** dialog; this dialog is accessible from the global Message Analyzer **Tools** menu. At any time, you can specify the default data viewer of your choice for all sessions, as described in Setting the Default Session Viewer, although you still have the option to choose a different one prior to starting a new session. However, the default data viewer setting will persist in each subsequent reopening of the **New Session** dialog, which provides you with the same option again to use the default viewer or choose another one. In addition, after session results display in the viewer you specified in session configuration, you have the option to display data in any of the other available data viewers to enhance data analysis by selecting them from the **New Viewer** drop-down list, as described in Locating Data Viewers for Selection. The data viewers that are available in Message Analyzer are described in the Data Viewers section.

> **NOTE**
> If you specify additional data viewers for the results of a particular Data Retrieval Session or Live Trace Session, Message Analyzer will automatically repopulate the original data set to the additional viewers in the same session.

Locating Data Viewers for Selection

You can specify any of the available Message Analyzer data viewers from the common data viewer Library that is accessible in each of the following locations:

- **Data Retrieval Session** configuration — provides data viewer selections from the common **Start With** drop-down list in the **New Session** dialog during session configuration.
- **Live Trace Session** configuration — provides data viewer selections from the common **Start With** drop-down list in the **New Session** dialog during session configuration.
- **New Viewer** — a drop-down list that is accessible from the following locations during an Analysis Session:
  - Global Message Analyzer **Session** menu, but only when a **Chart** is currently being displayed.
  - Global Message Analyzer toolbar.
  - The **Session Explorer Tool Window** context menu that is accessible by right-clicking any session node or viewer node in the **Session Explorer** window. The context menu provides access to the

New Viewer drop-down list, from where you can choose either a data viewer or a **Chart** viewer **Layout**.

The **New Viewer** drop-down list provides the same data viewer selections that are accessible from all the previously described locations. The different data viewers that you specify from the **New Viewer** drop-down list work with Live Trace Session or Data Retrieval Session *results* only, as previously indicated.

Specifying a different data viewer typically adds a separate viewer tab that contains an alternate data presentation format for the specific session that you choose, while retaining all other data viewers already in place for the chosen session. Note that the different data viewers that you specify for the *results* of a Live Trace Session or Data Retrieval Session create various data presentations, but do not modify the original contents of the trace or loaded message collection.

Setting the Default Session Viewer

You have the option to set a default data viewer for all sessions by selecting one from the **Default Viewer** drop-down list in the **Default Profile** pane on the **Profiles** tab of the global **Options** dialog. You can access this dialog from the global Message Analyzer **Tools** menu. Message Analyzer ships with the **Analysis Grid** as the default viewer for all sessions, but you can change to another viewer to customize your data analysis environment, as follows:

- **Global mode** — you can change the default data viewer that will be used by all sessions, from the **Analysis Grid** factory setting to any viewer that you choose in the **Default Viewer** drop-down menu in the **Default Profile** pane, as previously described.
- **Session specific mode** — you have the option to use a data viewer of choice whenever you start a session, regardless of what the default setting is, by clicking the **Start With** drop-down list in the **New Session** dialog and selecting a viewer.

> **NOTE**
> Because the **Details**, **Field Data**, **Message Data**, **Message Stack**, **Decryption**, **Diagnostics**, and other tool windows — accessible from the **Windows** submenu of the global Message Analyzer **Tools** menu — are either message-specific windows (respond to message selection) or session-specific (snap-to data viewer selection) windows, they depend upon and interact with the data viewers. As a result, you cannot specify one of these windows to start a Live Trace Session or Data Retrieval Session.

> **TIP**
> When you have session results data for which you have specified multiple data viewers, for example, the **Grouping** and the **Analysis Grid** viewers, you can bring the message data contained in a particular viewer into focus by clicking its associated node in the **Session Explorer** window, or by selecting its associated viewer tab.

More Information

To learn more about Message Analyzer data windows that interact with data viewers, see the Tool Windows topic.

To learn more about how to configure and save custom **Chart** viewer **Layouts**, see Configuring Chart Viewer Layouts.

Common Data Viewer Features

2 minutes to read

This section describes message processing features that enable you to manipulate the data presentation format, context, and content across multiple data viewers that are common to any particular Analysis Session. The primary data viewer where you apply the data manipulation features described in this section is the default **Analysis Grid** viewer, but you can also apply them to other data viewers, for example, the **Chart**, **Grouping**, and **Gantt** viewers. However, this section describes only the data manipulation features as they apply to the **Analysis Grid** viewer, given that this is the most common data viewer. Moreover, you can apply these data manipulation features to other viewers in the same or different session and they will have a similar impact. But note in general that the application of the Message Analyzer data manipulation assets or other tools described in this section will impact only the data viewer that is currently in focus.

The common data viewer features are described in the following topics:

Using the Filtering Toolbar
Filtering Data Sources
Setting Time Shifts
Configuring Time Format Settings
Using and Managing Message Analyzer Aliases
Configuring and Managing Message Analyzer Unions

Using the Filtering Toolbar

6 minutes to read

Message Analyzer now conveniently integrates several high-profile filtering functions into a single user interface (UI) feature known as the Filtering Toolbar, which is shown in the figure that follows. Whenever Message Analyzer displays session results in any particular data viewer, the Filtering Toolbar appears by default below the viewer tab along with any other toolbar that displays by default for the specific data viewer, for example, the toolbars that appear with the **Analysis Grid** and **Gantt** viewers. Note that a separate instance of the Filtering Toolbar also appears in the **Grouping** viewer whenever you display it.

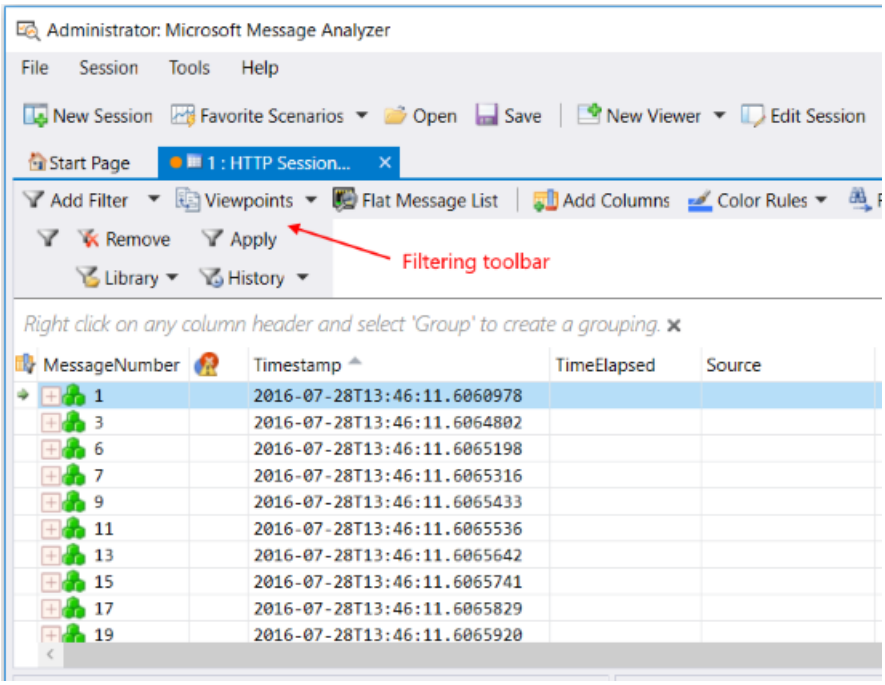

Figure 50: Message Analyzer Filtering toolbar

From the Filtering Toolbar, you can filter the trace results that are displayed by any in-focus data viewer, without affecting data in any other viewer or session. Moreover, by applying filtering to message data that is displayed in one viewer, there is no impact on the data displayed in any other viewer in the current session or data that is displayed in any viewer in any other session. Because you can apply filtering independently to each Message Analyzer data viewer, you gain more control over how you modify the data display for analysis.

Consider that you might want to view data in different types of viewers with a different filter applied to each, in order to create a unique analytical focus on each set of filtered data. You might also take things to the next level by creating a *tiered* configuration of multiple filters of different types or with a different focus for a particular data viewer, and then selectively apply them to — or remove them from — the trace results, independently of interaction with any other data viewer. This enables you to present different perspectives on filtered data all within the context of a single data viewer alone, as described in Working with Tiered Filtering Configurations.

Working with the Filtering Toolbar Features

To facilitate the previously described capabilities, Message Analyzer provides three main filtering functions on the Filtering Toolbar, as follows:

- **Add Filter** — each time you click this drop-down control, a new **Filter** panel appears by default along with the controls you need to specify and apply a view **Filter** to a set of trace results displayed in the current viewer, or to remove it from those trace results, as described immediately below. You can select

any of three different types of filters from this drop-down list, as follows:

- **Add Filter** — click this list item to display the **Filter** panel that contains a text-based editing surface for Filter Expressions and other controls you need for working with view **Filters** that you apply to the viewer that is currently in-focus. The controls include a **Library** from which to select a built-in view **Filter**, an **Apply** and **Remove** button to apply and remove the action of a **Filter**, and a **History** drop-down list that maintains a list of the last ten **Filters** that you applied. Note that a single **Filter** panel with associated controls appears by default whenever Message Analyzer displays session results in a data viewer. However, you can add as many more Filter panels as you need for configuration and selective application of specific filters.

  When you select a **Filter** from the **Library**, the **Filter** code automatically displays in the Filter Expression text box. Note that you can also compose your own **Filters** in this text box, but you may need to learn more about the Message Analyzer filtering language to do so, as described in Writing Filter Expressions.

> **IMPORTANT**
> To observe the filtering action associated with any view **Filter** that is configured in the Filter Expression text box, you must click the **Apply** button.

> **NOTE**
> The built-in view **Filters** of the centralized Filter Expression **Library** are available from the **Message Analyzer Filters** asset collection, which installs by default with Message Analyzer. You can **Manage** items in this collection from the **Library** drop-down list while other management features are available from the **Asset Manager** dialog that you can open from the global Message Analyzer **Tools** menu.

More Information
To learn more about working with view **Filters**, see Applying and Managing Filters.

- **Add Time Filter** — click this list item to display the **Time Filter** panel that contains the time window slider controls along with an **Apply** and **Remove** button. After you configure a window of time in which you want to view data, click the **Apply** button to start the filtering action. Thereafter, you can click the **Remove** button to disable the **Time Filter** effects and return to the original set of trace results, or you can configure and apply another window of time in which to view data.

> **TIP**
> You have the option to manually specify time stamp values in the **Time Filter** panel **Start Time** and **End Time** text boxes.

More Information
To learn more about working with **Time Filters**, see [Applying a Time Filter to Session Results] (applying-a-time-filter-to-session-results.md).

- **Add Viewpoint Filter** — this drop-down item is enabled only after you have applied a **Viewpoint**, as described below, otherwise it is disabled. After you apply a **Viewpoint**, which removes all messages above the selected **Viewpoint** for focused analysis, you might want to drill down further into the data that is displayed at the chosen **Viewpoint**. To achieve this, click the enabled **Add Viewpoint Filter** list item to display the **Viewpoint Filter** panel, from where you can configure and apply a **Viewpoint Filter**.

Note that this panel has the same controls that exist on the **Filter** panel. The only difference between applying a

Viewpoint Filter and a view **Filter** is that the action of the former works within the context of the current message set resulting from **Viewpoint** application, while the latter applies to all messages in the current set of overall trace results.

More Information
To learn more about **Viewpoint Filters**, see Applying Viewpoint Filters .

- **Viewpoints** — click this drop-down list to display the available **Viewpoints** that you can select and apply to a set of trace results. A **Viewpoint** enables you to look at your trace data temporarily from the perspective of the protocol, module, or layer represented by the selected **Viewpoint**, with no messages above it. Moreover, a **Viewpoint** drives the messages of a particular protocol, module, or layer to top-level to create a focused view of the types of messages you want to analyze, for example, SMB, TCP, Ethernet, and so on.

  A special **Viewpoint** that you can apply is **Disable Operations**, which changes the way top-level message nodes are organized and displayed in the **Analysis Grid**. When you apply this **Viewpoint**, Operations are broken apart and their constituent messages are placed in their original chronological sequence, to provide an alternate perspective for data analysis.

More Information
To learn more about **Viewpoints**, see Applying and Managing Viewpoints.
To learn more about **Operations**, see Working With Operations.

- **Flat Message List** — click this button on the Filtering Toolbar to create a message display that is similar to the Network Monitor view, where Operations are removed and messages, including fragments, are reorganized into their original chronological sequence.

More Information
To learn more about flattening the message display, see Creating a Flat Message List.

The following topics in this section provide further details on how to use the different components of the Filtering Toolbar:

Applying and Managing Filters
Applying a Time Filter to Session Results
Applying and Managing Viewpoints
Working With Operations
Creating a Flat Message List

See Also

Working with Tiered Filtering Configurations
Writing Filter Expressions
Grouping Viewer

Applying and Managing Filters

19 minutes to read

After you display message data in one or more of the Message Analyzer data viewers, you can apply a view **Filter** to reduce the scope of the data presented in a viewer according to filtering criteria that you define. This enables you to create a concise focus on the data you want to analyze. A view **Filter** enables you to target and isolate specific information for presentation and analysis, while preserving the original contents of your session results. For example, after you **Apply** a particular **Filter** from the Filter Expression **Library**, you can simply undo the filtering action by clicking the **Remove** command on the **Filter** panel, which appears whenever you click the **Add Filter** button on the Message Analyzer Filtering Toolbar. By executing the **Remove** command, Message Analyzer redisplays the trace results that existed immediately prior to applying the **Filter** you are removing.

Note that all built-in **Filters** are available from a common **Library** and are based on the Filtering Language that is described in Writing Filter Expressions. This **Library** contains the same Filter Expressions that are available as **Session Filters** when configuring a Live Trace Session or a Data Retrieval Session. The built-in **Filters** are provided by the **Message Analyzer Filters** asset collection, which is included with every Message Analyzer installation.

Generating Filters

You can display the controls and features you will need to create, apply, and remove one or more **Filters** by clicking the **Add Filter** button on the Message Analyzer Filtering Toolbar that appears above any in-focus session viewer tab. A single set of the indicated controls and features displays on a Message Analyzer **Filter** panel by default; however, you can display additional Filter panels for enhanced filtering, as described in Using the Filtering Toolbar. To generate the code for a view **Filter**, use any of the following methods:

- Select a built-in view **Filter** from the centralized Filter Expression **Library** drop-down list that appears whenever you click the **Add Filter** button on the Message Analyzer Filtering Toolbar.
- Utilize the IntelliSense statement completion service to write your own filters in any Filter Expression text box that appears when you click the **Add Filter** button on the Message Analyzer Filtering Toolbar.
- Automatically and quickly create view **Filter** code with the right-click context menu in the **Analysis Grid** viewer, as described in Creating Filters from the Analysis Grid Context Menu.
- Automatically and quickly create view **Filter** code with the right-click context menu in the **Details Tool Window**, as described in Creating Filters from the Details Tool Window Context Menu.

> **NOTE**
>
> When you use the methods described in the last two bullet points above to generate view **Filter** code, the Filter Expression code is added to the Filter Expression text box of the default **Filter** panel that appears whenever trace results are displayed in a data viewer.

After you generate the view **Filter** code, you must **Apply** the **Filter** for it to take effect, as described in Applying a Filter. Note that an **Apply** button appears in the default **Filter** panel and in each subsequent **Filter** panel that displays when you click the **Add Filter** button on the Message Analyzer Filtering Toolbar. Each **Filter** panel also contains the **Remove**, **Library**, and **History** controls.

Applying a Filter

By default, the filtering action of a view **Filter** impacts only the data viewer where you apply the view **Filter**, meaning that its action is specific to the current in-focus viewer only. The default action is initiated by clicking the **Apply** button on the **Filter** panel that is associated with the Filter Expression that you want to trigger.

> **TIP**
>
> You can also apply a view **Filter** by using the keyboard shortcut `Ctrl+Enter` and you can remove an applied view **Filter** by using the keyboard shortcut `Ctrl+Shift+Enter`. However, the Filter Expression text box in which the view **Filter** code displays must have the focus for this to work properly.

> **NOTE**
>
> A view **Filter** does not alter the original message data that you capture live or load into Message Analyzer. Whenever you run a Live Trace Session or Data Retrieval Session, a View Journal is automatically created as a repository for the results. A view **Filter** simply allows you to return a subset of View Journal data to your session viewer based on specified filtering criteria, for analysis purposes.

Working with Tiered Filtering Configurations

The Message Analyzer Filtering Toolbar contains the features that enable you to create, apply, and remove multiple view **Filters** against a set of trace results. With these features, you can create tiers of filtering that give you a finer level of control over the filtering process and how you display the data you want to examine. You can enhance your analysis process by selectively applying or removing any view **Filter** or combination thereof that exists in a set of **Filters** that you created. Consider that if you have a tiered configuration of two or more Filter Expressions, you can alternately select or unselect each view **Filter** to enable or disable it, respectively, and obtain different results based on different combinations of filtering criteria, for enhanced analysis.

For instance, if you create a view **Filter** such as `IPv4.Address==192.168.1.1` that enables you to isolate message conversations in which a specified address participated, you may want to then drill down further into the results to see whether a particular port carried the IP conversations. Instead of discarding the first **Filter**, replacing it with a new Filter Expression, and losing the current set of filtered results, you can click the **Add Filter** button on the Filtering Toolbar to create a new view **Filter** instance such as `UDP.SourcePort==500` to further isolate the data based on that criteria. Thereafter, you can alternately remove and reapply either of these view **Filters** to conveniently display a different set of results that enhances your analysis perspective, or you might **Remove** both view **Filters** to return to the original data set without losing any of the Filter Expression code that you specified.

Sample Tiered Filtering Configuration

An example of creating a tier of view **Filters** that you can apply and remove is shown in the figure that follows and is described in the procedure below.

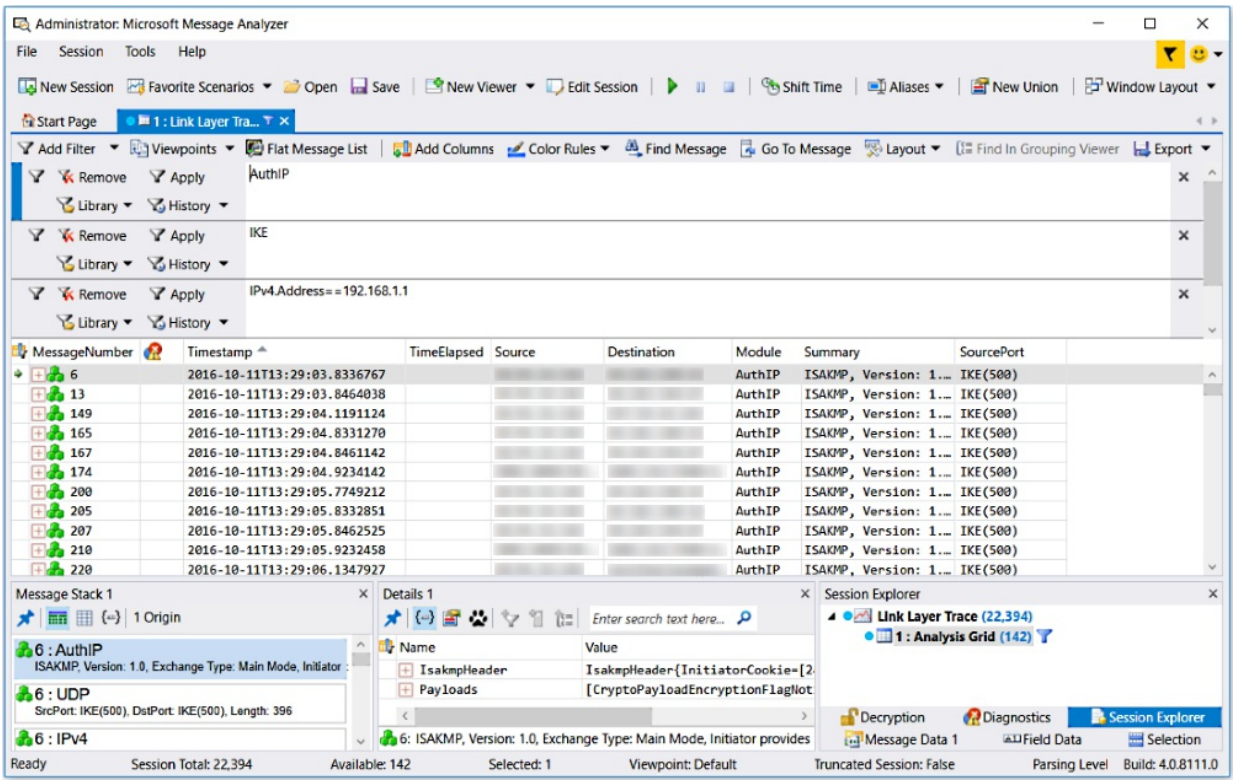

Figure 51: Tiered Filter Configuration with AuthIP Filter Applied

This example enables you to isolate authentication traffic from the Internet Key Exchange (IKE) cryptographic protocol and the Authenticated Internet Protocol (AuthIP), which is a Microsoft proprietary protocol that is an extension of IKE. Note that AuthIP provides a second authentication level to standard IKE authentication to add support for user-based authentication using Kerberos v5 or SSL certificates. The procedure captures data at the Link Layer with the **Microsoft-Windows-NDIS-PacketCapture** provider and makes use of four different view **Filters** that you can apply and remove in any combination. This simple example shows you how to dice and slice the data in various ways to obtain unique perspectives for analysis.

To create a tier of Filters that enhance the analysis process

1. From the **Start** menu, **Start** page, or task bar of a target computer running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system, click the **Microsoft Message Analyzer** icon to launch Message Analyzer. If you have not logged off and back on after first installing Message Analyzer, then start Message Analyzer with the right-click **Run as Administrator** option.

2. On the Message Analyzer **Start Page**, click the **Start Local Trace** button to begin capturing data at the Link Layer with the **Microsoft-Windows-NDIS-PacketCapture** provider.

   While Message Analyzer is accumulating messages in the default data viewer, typically the **Analysis Grid**, initiate any action that can invoke Kerberos authentication, such as file server resource access or some other site sign-in process. Note that this can occur automatically in Transport Layer Security (TLS) negotiations.

3. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.

4. Expand the **TCP** or **UDP** node in the **Field Chooser Tool Window** and navigate the message hierarchy until you find the **SourcePort** field; then double-click it to add it as a new column in the **Analysis Grid** viewer.

   This column is immediately populated with the port numbers used by the corresponding protocol or module messages in your trace results. Note that the display of column values will change as you apply specific view **Filters** in this procedure.

5. On the Message Analyzer Filtering Toolbar above the **Analysis Grid** viewer, click the **Add Filter** button

three times to display an additional 3 sets of **Filter** panels and Filter Expression text boxes in which to create view **Filter** code. Note that a single **Filter** panel displays by default whenever Message Analyzer displays session results.

6. In the first Filter Expression text box, type the code `AuthIP`, to create an atomic filter that returns messages from the AuthIP protocol while filtering out everything else, with exception of the AuthIP stacks.

   Note The meaning of an atomic filter is described in Creating Filters from the Analysis Grid Context Menu.

7. Click the **Apply** button on the Filter panel associated with the `AuthIP` filter to provide a concise set of AuthIP messages for analysis.

8. In the second Filter Expression text box, type the code `IKE`, to create an atomic filter that returns messages from the IKE protocol while filtering out everything else, with exception of the IKE stacks.

9. Click the **Remove** button on the **Filter** panel associated with the `AuthIP` filter and then click the **Apply** button on the **Filter** panel associated with the `IKE` filter to provide a concise set of IKE messages for analysis.

10. In the third Filter Expression text box, type the code `IPv4.Address==<192.168.1.1>`, while substituting appropriately for the placeholder address value in italics.

    When applied, this filter will return only the messages that contain either a **Source** or **Destination** address that matches the specified IP address, thereby providing a concise set of IP conversations where IKE negotiations took place.

> **NOTE**
> If you want to see only the AuthIP conversations, **Apply** the `AuthIP` filter and **Remove** the `IKE` filter.

11. In the fourth Filter Expression text box, type the code `UDP.SourcePort==500`.

    When applied, this filter will return only those AuthIP or IKE messages on UDP port 500 that provided authentication processes in the Internet Security Association and Key Management Protocol (ISAKMP) framework of authentication and key exchange.

12. Selectively apply and remove filters in the current tiered configuration in any combination, to expose different sets of information that can be useful for analysis and troubleshooting perspectives.

    For example, if you **Remove** the first two filters `AuthIP` and `IKE` and you **Apply** the second two filters `IPv4.Address==<192.168.1.1>` and `UDP.SourcePort==500`, Message Analyzer will show you all the AuthIP and IKE conversations that transited port 500.

> **NOTE**
> You are advised to experiment with different combinations of these view **Filters** so you can learn how to use tiered filtering configurations to expose different filtered results sets in a single data viewer and thereby enhance your data analysis process.

Using the Filter Expression Library

Message Analyzer provides a default set of built-in Filter Expression items that are accessible from the **Library** drop-down list that appears whenever you click the **Add Filter** button on the Filtering Toolbar. These **Filters** are sourced from the **Message Analyzer Filters** asset collection that you can manage from the **Manage**

Filter dialog, which displays when you click the **Manage** item in the indicated **Library** drop-down list. You can also share this collection or any part of it (including any **Filters** that you have created) with others, by using the **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu.

More Information
To learn more about the Asset Manager, see the Asset Manager topic.

The centralized Filter Expression **Library** contains built-in **Filters** that you can apply as view **Filters** to data displaying in a chosen data viewer, simply by selecting the **Filter** in the **Library** drop-down list on a displayed **Filter** panel. For example, you might apply the built-in Filter Expression `*SourcePort == IANA.Port.SMB` to the **Chart** viewer with a specified **Layout** to filter for messages of any protocol that have a **SourcePort** field equal to 445. You could then double-click some element in the **Chart** viewer **Layout**, for example, a bar element in the **Bar** visualizer component to which the filtering was applied, to automatically display the associated messages in a new instance of the **Analysis Grid** viewer for further examination.

You can also create your own Filter Expressions and add them to the **Library**, as described in Adding a Custom Filter to the Library. Any Filter Expressions that you create and add to the **Library** are stored under a subcategory that you specify under the top-level **My Items** category. However, note that the default set of built-in **Filters** are all contained in the top-level **Message Analyzer** category, with the exception of an **Example** filter that is placed in the **My Items** category by default, for use in Filter Expression development.

Compiling and Applying a Filter

As previously described, you can create your own custom Filter Expression to apply to a set of trace results that are displayed in a chosen data viewer. However, if you create your own Filter Expression, it is subject to successful compilation verification; otherwise you will be unable to use it. Note that Message Analyzer automatically performs a compilation verification of any Filter Expression that you specify on a particular **Filter** panel after you click the **Apply** button on the same panel. This ensures that you have a valid Filter Expression before it is applied to your trace results. If the Filter Expression does not pass the compilation check, an error message displays. At this point, you will need to correct the expression or abandon it. If the Filter Expression does pass the compilation check, the Message Analyzer Runtime will then apply the filter to your trace results.

> **NOTE**
> A similar compilation check is also applied to any **Session Filter** that you specify in the **New Session** dialog after you click the **Start** button in the dialog to begin a session, whether you are retrieving saved data or capturing live data.

Adding a Custom Filter to the Library

If you intend to add a custom-created **Filter** to the centralized Filter Expression **Library** for future use or to share with others, you will first need to display the **Edit Filter** dialog by selecting the **New Filter** item in the **Library** drop-down list on any **Filter** panel where you intend to configure the Filter Expression. From this dialog, you can specify **Name**, **Description**, and **Category** information. You can also write the code for the **Filter** in the Filter Expression text box of this dialog. However, if your custom filter code already existed in the Filter Expression text box of the **Filter** panel with which you are working at the time you launched the **Edit Filter** dialog, the Filter Expression code will have been automatically transferred to the dialog for your convenience. After you click the **Save** button in the dialog, Message Analyzer automatically performs a compilation check to ensure that the Filter Expression successfully compiles before saving it to the **Library** as a new asset. If the Filter Expression is invalid, a **Compile Query Error** message displays. Otherwise, you can assume that compilation succeeded.

TIP

You can expose the code for any built-in **Filter**, modify it, and then save it under a different name in the **My Items** category of your Filter Expression **Library**. To do this, select the **Create a Copy** command that displays in the context menu that appears when you right-click the **Filter** in the **Manage Filter** dialog, modify the Filter Expression code, and then save it with a new **Name** and in a specified **Category**.

Creating Filters from the Analysis Grid Context Menu

You can create and apply a view **Filter** very quickly to your data by right-clicking a data field value in most columns of the **Analysis Grid** viewer column layout and selecting the **Add '<columnName>' to Filter** command from the context menu that displays. The *columnName* value in the indicated command is a placeholder for the actual name of the **Analysis Grid** viewer column containing the data value that you right-click. The column name is automatically retrieved and displayed in the right-click menu, and when you select it, Message Analyzer builds a Filter Expression based on existing message field data values. For example, by clicking an IPv4 address in the **Destination** column, Message Analyzer automatically builds a Filter Expression such as `IPv4.Destination==192.168.1.1` and adds it to the Filter Expression text box that displays in the default **Filter** panel. Moreover, by clicking a **Module** column value such as **TCP**, Message Analyzer creates the atomic Filter Expression `TCP`. Note that as a result of the way these filters are created, they are guaranteed to return results.

NOTE

A Filter Expression such as `TCP` is called an *atomic* filter in Message Analyzer because it is a simple, left-hand-side-only filter that does not use an "equals" sign or any operators or combinators such as OR, AND, or NOT.

You can also save any Filter Expression that you created with the previously specified right-click method. You can also save Filter Expressions that you create in a similar manner from the **Details Tool Window**, as described ahead. For example, when you use the right-click method to create a Filter Expression, the text of the target Filter Expression is automatically transferred to the Filter Expression text box of the default **Filter** panel. To save this filter to your Filter Expression **Library**, first click the **Library** drop-down list and select the **New Filter** command to display the **Edit Filter** dialog. You can then specify a subcategory under the **My Items** top-level category of your **Library**, optionally add **Name** and **Description** information, and then click the **Save** button in the **Edit Filter** dialog to save the new Filter Expression.

Creating Filters from the Details Tool Window Context Menu

Similar to the way you create a view **Filter** from the **Analysis Grid** viewer context menu, you can also create a view **Filter** from the **Details Tool Window**, by right-clicking any field in the **Name** column of the **Details** window and selecting the **Add '<fieldName>' to Filter** context menu item. The *fieldName* value in this command is a placeholder for the actual field name in the **Name** column of the **Details** window.

TIP

If the **Details** window is not displayed, select the **Details 1** item from the **Details** drop-down list in the **Windows** submenu on the Message Analyzer global **Tools** menu to restore it.

As in the case of creating a view **Filter** from the **Analysis Grid** viewer, selecting the previously indicated context menu command in **Details** for creating a Filter Expression only adds its code to the Filter Expression text box of the default **Filter** panel. To see the results of the automatically configured view **Filter**, you must click the **Apply** button in the default **Filter** panel.

Managing Filters as Shared Items

Your local Filter Expression **Library** contains the default **Message Analyzer Filters** asset collection of **Filter** items plus any items that you create, and you can share all of these items with others. To do this, Message Analyzer provides a simple way to expose your Filter Expression items to others for sharing, or to retrieve Filter Expressions that others have shared. You can share your Filter Expression **Library** items directly with others by using the **Export** feature in the **Manage Filter** dialog to save one or more Filter Expression items to a designated file share. You can also use the **Import** feature in the same dialog to access Filter Expression items that have been shared by others. The **Manage Filter** dialog is accessible by selecting the **Manage Filters** item from the **Library** drop-down list on the **Filter** panel with which you are working.

Sharing Filters on a Feed

You can share your Filter Expression items through a user feed that you configure in the Message Analyzer Sharing Infrastructure. You can create your own feed from the **Settings** tab of the Message Analyzer **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu, and the feed (defined by a designated share or other location) will appear on the **Downloads** page of the same dialog. Thereafter, you can update existing Filter Expression items or add others and make them available to team members or other users through the configured feed, where they can view, synchronize, and download them from the **Downloads** or **Settings** tabs of the **Asset Manager** dialog. However, the synchronization aspect of the publishing feature on user feeds requires some manual configuration at this time to enable updates, as described in Manual Item Update Synchronization.

Receiving Filter Asset Collection Updates

Message Analyzer also has a default **Message Analyzer** subscriber feed on the **Downloads** tab of the **Asset Manager** dialog that enables you to download the **Message Analyzer Filters** asset collection from a Microsoft web service and to synchronize with asset collection updates that are periodically pushed out by the service, as useful Filter Expressions are developed at Microsoft for the community of Message Analyzer users. To receive these updates that will appear in the **Message Analyzer** category of your local Filter Expression **Library**, the **Message Analyzer Filters** asset collection should be in the auto-sync state (circular icon with up and down arrows) on the **Asset Manager** dialog **Settings** tab. At any time, you can perform a download of an auto-synced collection from the **Settings** tab of the **Asset Manager**.

> **IMPORTANT**
>
> The **Message Analyzer Filters** collection is installed by default with Message Analyzer, so it is unnecessary to download the collection each time you start Message Analyzer. But if it is the first time you have started Message Analyzer, you are presented with a **Welcome** dialog that provides you with the choice to opt in or out of automatic updates. If you choose to opt in to auto-syncing updates, then all Message Analyzer asset collections are automatically set to the auto-sync state, including the **Message Analyzer Filters** collection, and no further action is required.

However, if you opted out, you still have the option to automatically receive periodic collection updates later by setting the **Offline** mode to **Online** on the **Downloads** tab of the **Asset Manager** dialog and clicking the **Sync All Displayed Items** button. This action auto-syncs all asset collections; however, you can set individual collections to the auto-sync state on the **Downloads** tab as you require them. To do this, click the download icon to the right of the collection on the **Downloads** tab and select the **Automatically sync item collection updates when available** option in the **Item Download Options** dialog.

More Information

To learn more about applying view **Filters**, see Filtering Message Data.
To learn more about the Filtering Language and how to write filter expressions, see Writing Filter Expressions.
To learn more about sharing Message Analyzer Library items, including further details about the common

Manage <_AssetType_> dialog, see the Sharing Infrastructure topic.
To learn more about auto-syncing item collections, see Managing Asset Collection Downloads and Updates.

See Also

Using the Filtering Toolbar
Applying and Managing Viewpoints
Applying a Time Filter to Session Results

Applying a Time Filter to Session Results

6 minutes to read

Message Analyzer enables you to filter messages in data viewers that you have opened, such as the **Analysis Grid** and **Chart** viewers, based on a configured time window in which you elect to view messages. This feature is called a **Time Filter** and it is accessible from the **Add Filter** drop-down list on the Message Analyzer Filtering Toolbar that displays above the analysis surface where data viewers appear when selected. You can configure a **Time Filter** for Live Trace Session results similar to the way you configure a **Time Filter** for a Data Retrieval Session with some exceptions, as described in Applying an Input Time Filter to a Data Retrieval Session, which includes employing controls that you can use to define the **Start Time** and **End Time** of a time window in which to view messages. However, the scope of these **Time Filters** is different, as are some label and naming nomenclatures of the user interfaces involved.

The difference in scope of these **Time Filters** has to do with the context in which they are applied. For example, if you apply a **Time Filter** to data that you load into Message Analyzer through a Data Retrieval Session, only the messages that fall within the time window that you specify in the time-filter configuration are retrieved. All other messages are filtered out when the data is loaded, which therefore makes such filtered messages unavailable in the chosen Message Analyzer data viewer. Thereafter, the filtered-out messages can be restored only by undoing the **Time Filter** configuration and reloading the data from the **Edit Session** dialog. On the other hand, although a **Time Filter** for a set of trace results applies the same type of time-slot filtering, it does so in the context of an already parsed message collection. This means that the original raw trace data is preserved in the Message Store even if the message display is altered by the way you manipulate the data. As a result, a **Time Filter** for trace results provides the convenience of enabling you to toggle back and forth between the original trace results and the time-filtered trace results, which can be a benefit to your analytical processes.

Considering Performance vs. Usability Factors for Time Filter Application

You can apply a **Time Filter** against a set of trace results after it is loaded into Message Analyzer or you can apply an *input* **Time Filter** at the moment when you actually import the same data into Message Analyzer. To decide which is best in your circumstances, you might want to consider the tradeoffs between performance and usability, especially when loading data from very large files.

Applying a Time Filter When Retrieving Data

When you apply a **Time Filter** to a Data Retrieval Session, the amount of data being loaded is reduced to a specified time window, which results in less impact on CPU and memory resources, and therefore demonstrates better performance. However, the resulting data set has all messages removed that are outside the specified time window, which could have an impact on usability and data analysis in terms of message context and other relationships. Moreover, an input **Time Filter** might obscure conversation contexts in the filtered message collection that you import, whereas a **Time Filter** applied to a set of trace *results* does not. To avoid the loss of data in the case of the former, you would need to edit the Data Retrieval Session by setting a new **Time Filter** window in the **New Session** dialog and then reload the data. In the case of the latter, Message Analyzer enables you to remove or resize and reapply the **Time Filter** to the current set of trace results as necessary, making it more convenient to work with the data, given that a data reload is unnecessary.

Also note that for some input file types, Message Analyzer may not be able to determine the start and end times, in which case the **Time Filter** controls in the **New Session** dialog for a Data Retrieval Session will be disabled. However, you can always utilize the **Time Filter** feature after the data displays in your Analysis Session, as needed, regardless of the input file type.

More Information

To learn more about applying a **Time Filter** when loading data from logs and trace files into Message Analyzer, including the file types for which Message Analyzer can determine start and end times, see Applying an Input Time Filter to a Data Retrieval Session.

Applying a Time Filter to Live Trace Session Results

On the other hand, if you wait to apply a **Time Filter** to a set of messages that are already displayed in a data viewer such as the **Analysis Grid**, an applied **Time Filter** must check the time stamps on each message in the entire collection against the filtering criteria, which can impact performance if there is a very high message volume. However, usability is improved since you can easily restore the entire message set, including conversation contexts, with a single click of the **Apply** button in the **Time Filter** panel on the Filtering toolbar after altering the time window as needed, for ease of analysis.

Configuring a Time Filter for a Set of Trace Results

When you are ready to configure a **Time Filter** for a set of trace results, select the **Add Time Filter** command that displays when you click the **Add Filter** drop-down list on the Message Analyzer Filtering Toolbar that appears just above the analysis surface where all data viewers display. This action displays the **Time Filter** panel that contains **Start Time** and **End Time** text boxes that by default display the original time window boundaries of the displayed message collection, along with a set of time slider controls that enable you to adjust the window of time in which you want to view data.

As you change the time window with the slider controls, the time values in the **Start Time** and **End Time** text boxes display also change. When you finish adjusting the time slider controls, the corresponding new start and end time values define the selected time window in which you are choosing to view data.

> **TIP**
>
> You have the option to apply a **Data Source** filter to a set of trace results to isolate the messages from specified Data Source providers that were used during data capture, as described in Filtering Data Sources. By combining a **Data Source** filter with a **Time Filter**, you can obtain a very narrowly focused set of results that drills down to a specific window of time in which data was captured by a specific Data Source.

As you set the slider controls, Message Analyzer creates a Filter Expression in the background that will facilitate the time-filtering action. Note that you can manually specify time stamps in the **Start Time** and **End Time** text boxes; however, any time stamp that is outside the message collection time boundaries will create an erroneous **Time Filter**.

Applying a Time Filter

After you set the time window configuration for a **Time Filter** as previously described, you can apply the filter by clicking the **Apply** button on the **Time Filter** panel. Message Analyzer responds by displaying only the messages that fall within the time window you specified. If you want to return to the unfiltered display of original message data, click the **Remove** button on the **Time Filter** panel. This action removes the effects of the previously applied **Time Filter**, but does not alter the time window configuration that you specified. If you want to reapply the **Time Filter**, you can do so by clicking the **Apply** button again on the **Time Filter** panel. You can toggle back and forth between the filtered and unfiltered view configurations as many times as you want, as the time window values that you set for a particular session viewer will persist until you change them.

> **NOTE**
>
> As you alternately apply and remove a **Time Filter** configuration, you can observe the activation of session progress indicators in the **Session Explorer Tool Window**.

See Also

Using the Filtering Toolbar

Applying and Managing Viewpoints

16 minutes to read

In Message Analyzer, the default **Analysis Grid** viewer focuses on top-level messages to provide a compact display of data summaries, so that you can very quickly understand issues at a high level. As a result, other important details can be hidden in this view, such as the underlying origins messages that support Operations or other top-level messages. However, because awareness of the activities of specific protocols at the lower layers can be crucial to data analysis, it is often necessary to achieve a focused analytical perspective at these levels. Moreover, it can also be advantageous to be able to view only the traffic of a particular higher-layer protocol, for example, HTTP. In Message Analyzer, this is made possible through the application of **Viewpoints**.

Understanding Viewpoints

To make your troubleshooting efforts easier, Message Analyzer enables you to examine network traffic from the perspective of a protocol, where you can display the specific protocol messages at top-level in the **Analysis Grid** viewer with no layers above them. For this reason, **Viewpoints** could be considered *layer filters* because they temporarily remove the display of all messages above the applied protocol **Viewpoint**, such that only those protocol messages appear at top-level in the **Analysis Grid** viewer.

For example, when viewing trace results in the **Analysis Grid** viewer, you may have higher-layer traffic that obscures the underlying messages that you want to troubleshoot. By default, the **Analysis Grid** viewer displays top-level messages and Operations in single rows with expandable nodes, where the message origins or underlying message stack is concealed under multiple lower-level expansion nodes. As a result, you can only examine the details of messages in the underlying layers by expanding the message nodes one by one to expose the protocol or module layer that you want to troubleshoot. Repeating this process across an entire trace when searching for specific data can become extremely labor intensive, particularly in a large trace with many messages.

To alleviate this difficulty, Message Analyzer provides a set of pre-configured **Viewpoints** that enable you to expose the data for specific message types in top-level rows of the **Analysis Grid** viewer, with all the upper-layer messages above the **Viewpoint** level removed. However, even though the upper-layer messages are removed by a **Viewpoint**, you still have the option to view the message stack in its entirety for any viewpoint message — while the **Viewpoint** is applied — by opening the **Message Stack Tool Window**. Thereafter, when you select any viewpoint message in the **Analysis Grid** viewer, the selected viewpoint message and the layers above and below it display in the **Message Stack** window to give you some layering context, along with a quick view of **Summary** statistics for messages at each layer.

Application and Removal of Viewpoints

When messages are parsed by Message Analyzer, they are indexed. When you apply a **Viewpoint** to a set of parsed messages, Message Analyzer simply reorganizes the data display by retrieving messages whose indexes correlate with the applied **Viewpoint** filtering criteria. The result is that you can display the viewpoint messages at the top-most level in the **Analysis Grid** viewer, which can include all operations that exist at the current **Viewpoint**, if they exist at that level. For example, if you apply the **SMB/SMB2 Viewpoint**, then operations for the SMB and SMB2 protocol will display. The current exception to this is if an upper-layer protocol that is above a set **Viewpoint** also defines operations. In this case, operations for the latter protocol will display at top-level.

The **Viewpoint** that displays by default in the **Analysis Grid** viewer is a summary view of top-level messages that have no other message layers above them. After applying a **Viewpoint** to a set of messages and changing the data to the perspective of a particular protocol, you can return to the default **Viewpoint** by clicking the **No Viewpoint** item in the **Viewpoint** drop-down list on the Filtering Toolbar. You can also remove the current

Viewpoint by selecting another one from the **Viewpoint** drop-down list.

Accessing the Built-In Viewpoints

To assist your troubleshooting efforts, Message Analyzer provides a robust set of built-in **Viewpoints** that you can access and apply from the Filtering Toolbar that appears whenever Message Analyzer displays a set of trace results, as described in Using the Filtering Toolbar. These **Viewpoints** are contained in the **Message Analyzer Viewpoints** asset collection Library that is accessible from the **Viewpoints** drop-down list on the Filtering Toolbar. They enable you to filter and reorganize your data view to display messages from the perspective of different protocols, modules, or message layers, in accordance with the functionality of the applied **Viewpoint** . The application of a **Viewpoint** enables you to achieve a unique analytical perspective on data that might normally be hidden from view or difficult to expose.

Using the Viewpoint Features

The remaining topics in this section describe the built-in **Viewpoints** that ship with Message Analyzer, applying a **Viewpoint**, how to work with **Viewpoint Filters**, and using the features for managing **Viewpoints**:

Applying a Built-In Viewpoint — as applied to messages displaying in a data viewer.
Applying Viewpoint Filters — to enhance the analysis context.
Managing Viewpoints as Shared Items — to enable the mutual sharing of **Viewpoints** with others.
Receiving Viewpoint Asset Collection Updates from Microsoft — to synchronize your **Viewpoints** asset collection for automatic updates from Microsoft.

Applying a Built-In Viewpoint

By default, Message Analyzer provides numerous built-in **Viewpoints** that are described below. You can apply these to any set of messages that are displayed in any data viewer by selecting a chosen **Viewpoint** from the **Viewpoints** drop-down list on the Filtering Toolbar while a particular data viewer has focus.

> **TIP**
> After you apply a **Viewpoint** to such a data viewer, you can hover over the data viewer session tab or the viewer node in **Session Explorer** at any time thereafter to view a popup that indicates which **Viewpoint** is currently applied to the data set, along with any **Viewpoint Filter**, **Filter**, or **Message Range** filter (a **Bookmarks** or **Gantt** viewer context menu item) that is currently applied to the viewer.

The built-in **Viewpoints** that are contained in the **Message Analyzer Viewpoints** asset collection consist of the following:

- **No Viewpoint** — enables you to return to the original set of trace results where no **Viewpoint** is applied.
- **Data Link** layer — enables you to display messages at top-level from protocols related to the Data Link layer, such as the Ethernet, PPP, ARP, and WiFi protocols, and their origins.
- **Disable Operations** — enables you to disable Operations by initiating a reparse of your trace results that disables the Message Analyzer Runtime's default encapsulation of request and response message pairs in Operation nodes. This feature currently supports de-encapsulation of Operations for the HTTP protocol only.

  This command can be useful for filtering under an applied **Viewpoint**, when you can only use a **Viewpoint Filter** to isolate particular types of messages, for example, an HTTP Request or Response message. Otherwise, without Operations disabled, you would be unable to isolate such messages from the default Operation node encapsulation.

More Information
To learn more about Operations, see Working With Operations.

- **Transport Layers UDP/TCP** — enables you to display messages at top-level from the TCP and UDP protocols only, including their origin stacks.
- **TCP** — this **Viewpoint** reorganizes your data to enable easier diagnosis of the TCP layer. It places TCP messages on top, which can facilitate diagnosis of TCP performance issues that include the analysis of TCP SequenceNumber and AcknowledgementNumber values, TCP flags such as SYNs, SACKS, and ACKs, retransmits, broken three-way handshakes, window size, TCP options, and so on. Note that if you apply this **Viewpoint**, Operations will no longer be visible, as they typically exist at a layer above this **Viewpoint**.

> **TIP**
> To enhance your analytical perspective with the **TCP Viewpoint**, you can use the **TCP Deep Packet Analysis with Relative Sequence Number with Grouping** view **Layout** for the **Analysis Grid** viewer to display the relevant field data in a predefined Grouped configuration with **Network**, **Transport**, and **SourcePort** groups.

- **UDP** — provides perspective from the **Viewpoint** of the UDP transport protocol.
- **Network** layer — enables you to display messages at top-level from the IPv4, IPv6, DHCPv4, DHCPv6, and DNS protocols only, including their origins messages.
- **Ethernet** layer — enables you to display Ethernet messages at top-level with no further parsing.
- **ETW** — enables you to remove all messages above the ETW layer to expose and simplify event diagnostics. This **Viewpoint** can also make event analysis easier when you are developing message providers or other components that write ETW events.
- **HTTP** — an application-layer **Viewpoint** that places HTTP messages at top-level in the **Analysis Grid** viewer. Provides a convenient way to analyze the request/response pairs of HTTP Operations without having to search for the response messages. Also facilitates improved filtering for request and response messages, as described in Disabling Operations.

> **NOTE**
> It is possible that HTTP messages can be hidden within SOAP message stacks. If you apply the **HTTP Viewpoint** when this is the case, SOAP messages should disappear and HTTP messages will display at top-level. However, the HTTP messages may not display as Operations in this case.

- **SMB/SMB2** — an application-layer **Viewpoint** that places SMB and SMB2 messages at top-level in the **Analysis Grid** viewer by removing RPC and any other message layers on top, for example, GSSAPI and Kerberos messages.
- **SMB/SMB2 No Operations** — an application-layer **Viewpoint** that is identical to the **SMB/SMB2 Viewpoint**, except that Message Analyzer does not display any SMB/SMB2 Operation nodes in this context. Enables you to view SMB/SMB2 request and response messages in their original chronological order. Also facilitates improved filtering for request and response messages, as described in Disabling Operations.
- **WinInet (HTTP/s)** event layer — enables you to display and diagnose HTTP and unencrypted HTTPS events.
- **SOAP** — enables you to display messages at top-level from the SOAP protocol only, plus the origins

messages.

Applying Viewpoint Filters

Message Analyzer provides a Filter Expression **Library** on the **Viewpoint Filter** panel that appears when you click the **Add Viewpoint Filter** item in the **Add Filter** drop-down list on the Filtering Toolbar. This **Library** is the same centralized **Library** that is located on the Filter panel that appears when you click the **Add Filter** drop-down on the Filtering Toolbar. Because it is the same **Library**, the same Filter Expressions from the **Message Analyzer Filters** asset collection are available for selection as **Viewpoint Filters**. You also have the option to create and apply custom Filter Expressions of your own design, as described in Applying and Managing Filters.

> **NOTE**
> The controls on the **Viewpoint Filter** panel are described in Using the Filtering Toolbar.

Drilling Down to Expose Target Messages

From the **Viewpoint Filter** panel, you can **Apply** such filters to a set of messages that is already filtered by the criteria of an applied **Viewpoint**. The advantage of using a **Viewpoint Filter** is that it enables you to drill down further to expose messages of interest based on the additionally applied filtering. Obviously, the filtering you apply should be relevant to the **Viewpoint** context in which you are working. In a typical usage scenario, you might have already applied a view **Filter** to a set of trace results when you realize that you should set the **Viewpoint** to a particular layer so you can focus on a condensed and more relevant message set for your current analysis. After you select a chosen **Viewpoint**, all messages above the **Viewpoint** level disappear. Once the **Viewpoint** is set, you can drill down even further to isolate a message or messages that meet the criteria of a **Viewpoint Filter** that you specify. Whether you select a built-in or custom-designed **Viewpoint Filter**, you must click the **Apply** button on the **Viewpoint Filter** panel to initiate the filtering action. To undo such filtering action, click the **Remove** button on the same panel.

Filtering Behaviors

The behavior of view **Filters** and **Viewpoint Filters** is similar within the different contexts in which they are applied. More specifically, view **Filter** behavior with respect to an entire set of trace results is similar to the way a **Viewpoint Filter** behaves when applied to a **Viewpoint** results set. The difference in the **Viewpoint** filtering scenario is that you are able to generate a more precise focus on specific messages of interest *within the context of the applied **Viewpoint***. In general, **Viewpoints** enable you to create focus by removing all messages above the **Viewpoint** protocol/s. However, by also applying a **Viewpoint Filter**, you can be even more selective of the messages you are exposing in the **Analysis Grid** or other viewer for analytical purposes. The scenario described below may help explain the difference between using view **Filters** and **Viewpoint Filters**.

Viewpoint Filtering Example

If any top-level message or one of its origins messages in a set of trace results matches the criteria of a view **Filter**, Message Analyzer returns that top-level message and its origins (stack) when the **Filter** is applied. If you then apply a **Viewpoint** while the same view **Filter** is applied, the **Viewpoint** will cause any of the filtered messages that match the **Viewpoint** criteria to appear at top-level in accordance with the applied **Viewpoint's** functionality. The following example illustrates the results that occur when you apply a view **Filter**, a **Viewpoint**, and a **Viewpoint Filter** to a set of trace results.

Example: If you apply the view **Filter** `TCP.SequenceNumber == 1648745597` to an original set of trace results, and a top-level message or one of its origin messages has a field value that matches that TCP sequence number, then that top-level operation (SMB2 message #19 and its origins in the figure immediately below) is isolated in the **Analysis Grid** viewer with all other messages removed from display. Note that the message stack has been expanded in this figure to show SMB2 request and response messages #19 and #27 that comprise the Operation.

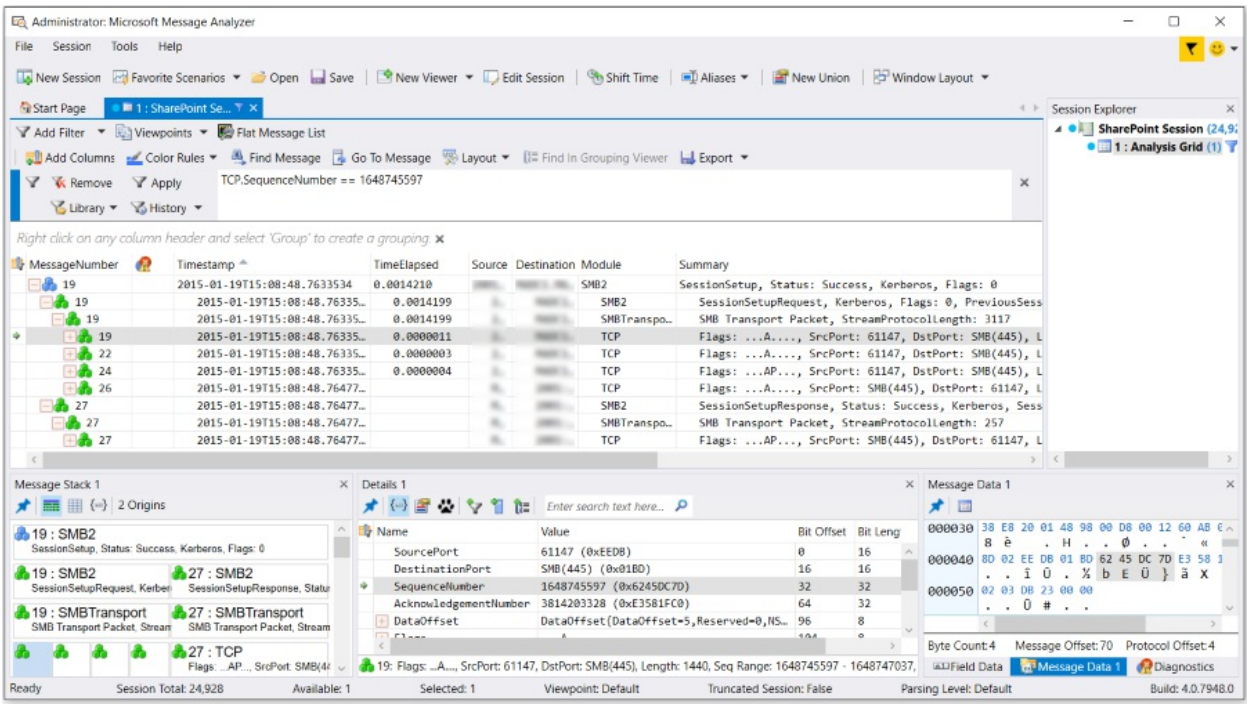

Figure 52: Message Analyzer TCP Filter Applied Only

The above result is expected behavior, as all view **Filters** work in this manner. If you now apply a **TCP Viewpoint** to this display configuration, all TCP fragments that are contained in the stack of the former top-level SMB2 message #19 are now pushed to top-level, so that only TCP messages display with nothing above that layer, which also includes the TCP fragment (#19) that met the filtering criteria of the applied view **Filter**, as shown in the figure below.

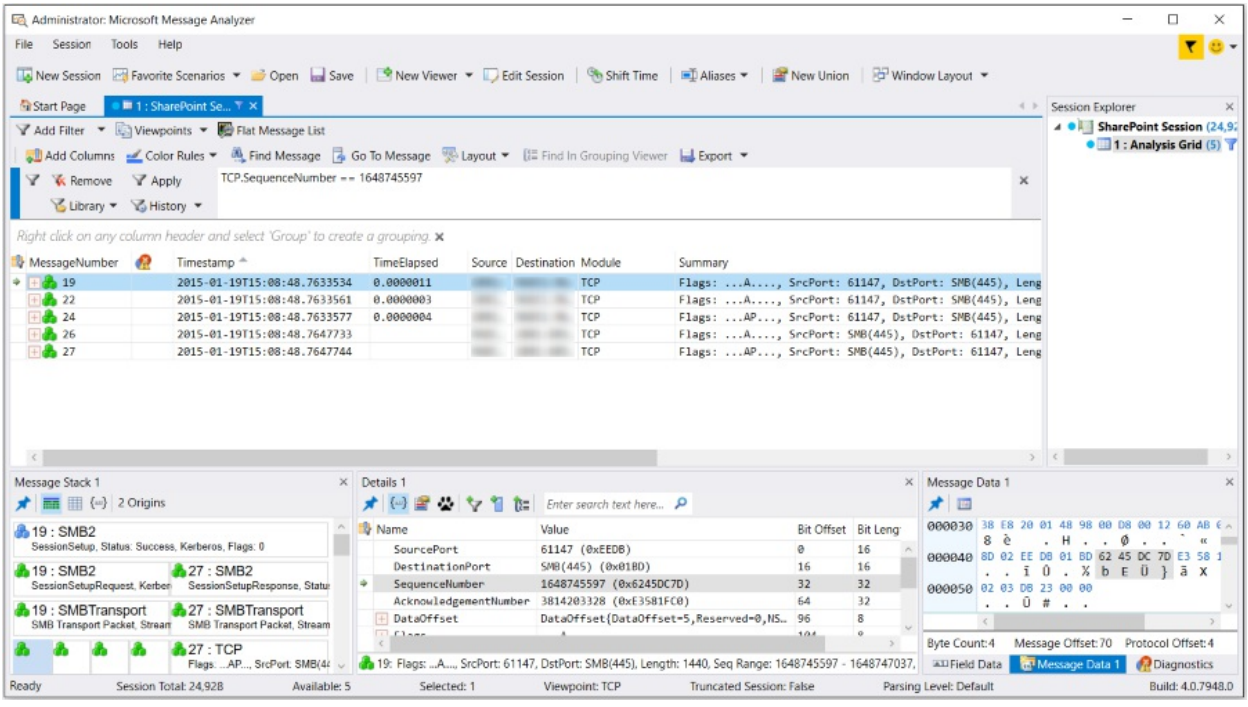

Figure 53: Message Analyzer TCP Filter and TCP Viewpoint Applied

The following figure shows the results after applying a **Viewpoint Filter** that is the same as the originally applied view **Filter** `TCP.SequenceNumber == 1648745597`. When you apply this same filter as a **Viewpoint Filter**, it results in displaying only the TCP fragment (#19) that meets that filtering criteria along with its underlying stack messages, as shown in the figure below. Optionally, you can remove the original view **Filter** before applying the indicated **Viewpoint Filter** to achieve the same result.

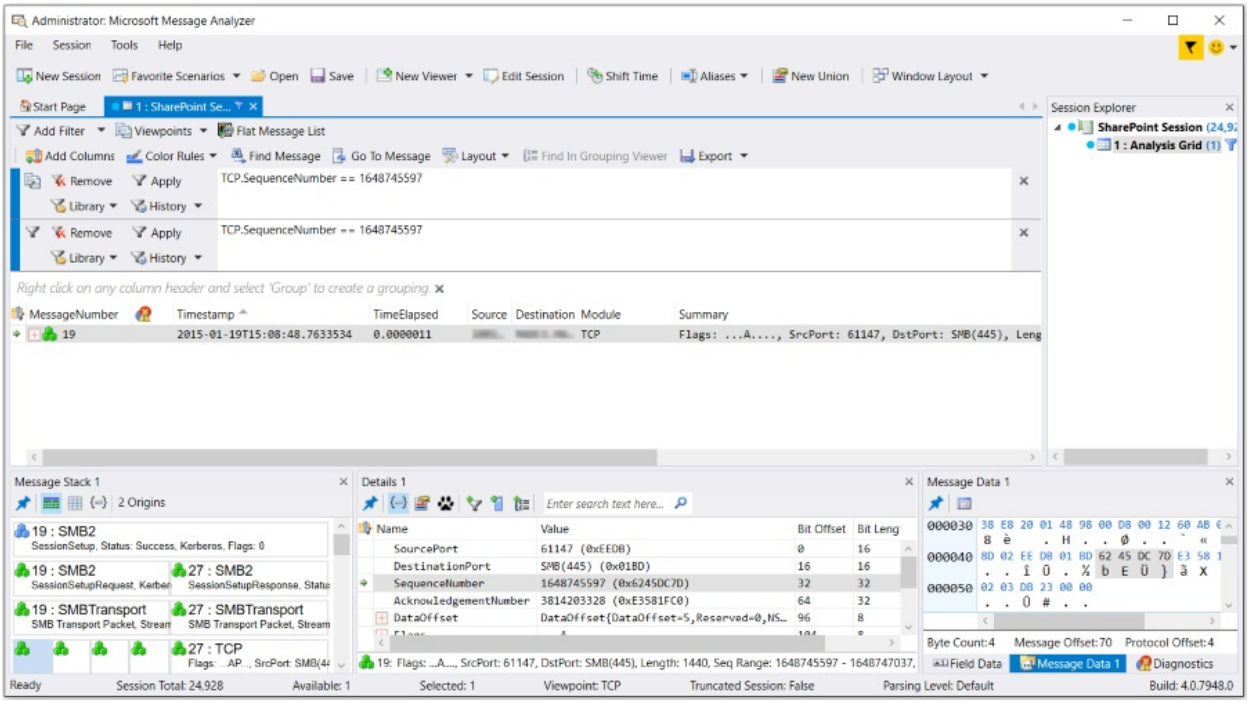

Figure 54: Message Analyzer TCP Filter, TCP Viewpoint, and Viewpoint Filter Applied

This behavior is similar to the way a view **Filter** works against an original set of trace results with no **Viewpoint** applied. In summary, the working differences between a view **Filter** and a **Viewpoint Filter** is the context in which the filtering is applied.

Managing Viewpoints as Shared Items

The **Viewpoint** items in the **Viewpoint** drop-down list on the Filtering Toolbar are shareable assets. Message Analyzer provides a simple way to expose these **Viewpoint** items to others for sharing purposes, or to retrieve **Viewpoint** items that others have shared.

> **NOTE**
>
> The ability to create new **Viewpoints** or edit existing ones may be available in a future Message Analyzer release.

Exporting a Viewpoint Asset Collection for Sharing

You can create and share a **Viewpoint** asset collection with others that consists of a full replica of the **Message Analyzer Viewpoints** asset collection or you can create and share a subset of the overall collection. To specify the items you want to include in the collection that you will post to a designated file share or other location, simply click the **Manage Viewpoints** item in the **Viewpoints** drop-down list on the Filtering Toolbar to open the **Manage Viewpoint** dialog. After the dialog opens, you can select items you want to include in the shareable collection by placing a check mark in the check box of each item you want to include in the collection. When complete, click the **Export** button on the toolbar of the **Manage Viewpoint** dialog to display the **Save Library** dialog, from where you can specify **Title**, **Description**, **Author**, and **Organization** information. When you click **OK** to exit this dialog, the **Select Library Location** dialog displays to enable you to navigate to a save location for the asset collection.

Importing a Shared Viewpoint Asset Collection

If you want to retrieve items from a collection that someone else has shared in the previously specified manner, click the **Import** button on the toolbar of the **Manage Viewpoints** dialog to open the **Select Library to Open** dialog. After you navigate to the location where the asset collection has been shared, click the **Open** button in the dialog. At this point, the **Select Items to Import** dialog displays, from where you can select the items in the collection that you want to import. You can also specify a category in which the collection items will appear. When complete, click the **OK** button, at which time the selected items are imported into your **Viewpoints** asset

collection.

Sharing a Viewpoint Collection on a User Feed
In addition, you can share a **Viewpoints** asset collection through a user feed that you configure in the Message Analyzer Sharing Infrastructure, which you can accomplish from the **Settings** tab of the **Asset Manager** dialog. This dialog is accessible from the global Message Analyzer **Tools** menu. Thereafter, you can use the **Export** feature of the **Manage Viewpoint** dialog to post your **Viewpoints** asset collection to the feed so that others can access them. Whenever you update the contents of this **Viewpoints** asset collection, you can make the changes available to team members or other users through the configured feed, where they can view, synchronize with, and download your asset collection items. However, to enable users to download asset collection updates, there is some manual configuration required at this time, as described in Manual Item Update Synchronization.

Receiving Viewpoint Asset Collection Updates from Microsoft

Microsoft provides a default **Message Analyzer** feed on the **Downloads** tab of the **Asset Manager** dialog that enables you to download the **Message Analyzer Viewpoints** asset collection from a Microsoft web service and to synchronize with asset collection updates that are periodically pushed out by the service.

The **Message Analyzer Viewpoints** asset collection is installed by default with Message Analyzer, so it is unnecessary to download the collection whenever you start Message Analyzer. But if it is the first time you have started Message Analyzer, you are presented with a **Welcome** dialog that provides you with the choice to opt in or out of automatic updates. If you choose to opt in to auto-syncing updates, then all Message Analyzer asset collections are automatically set to the auto-sync state, including the Message Analyzer **Viewpoints** asset collection, and no further action is required.

However, if you opted out, you still have the option to automatically receive periodic collection updates later by setting the **Offline** mode to **Online** on the **Downloads** tab of the **Asset Manager** dialog and then clicking the **Sync All Displayed Items** button to auto-sync all asset collections; or you can set individual collections to the auto-sync state on the **Downloads** tab as you require them. To do this, click the download icon to the right of the collection on the **Downloads** tab and select the **Automatically sync item collection updates when available** option in the **Item Download Options** dialog.

More Information
To learn more about the functions of the Filter Expressions in the **Message Analyzer Filters** asset collection, see Filtering Live Trace Session Results.
To learn more about the Sharing Infrastructure and managing user Library items, downloading asset collections, and auto-syncing asset collection updates, see the Sharing Infrastructure and Managing Asset Collection Downloads and Updates topics.

See Also

Using the Filtering Toolbar

Working With Operations

9 minutes to read

Message Analyzer provides some special features for working with protocols that employ request and response messages as the basis of their negotiation architecture. In other protocol analysis tools, it is a common practice to display messages in the order in which the tool originally captured them. In practice, this can present some difficulties for those who are analyzing such data, because they may have to search through literally hundreds, if not thousands of messages to locate the response to an associated request message. Although other techniques such as filtering can be used by the analyst to find these messages, it can still involve multiple steps of additional manual configuration to locate and organize the data before analysis can begin.

Message Analyzer provides a simple solution to this problem by encapsulating each request and response message pair of any protocol that uses this architecture into a single, expandable Operation node that displays as a top-level message row in the **Analysis Grid** viewer. By expanding such an Operation node, you can expose the original request and response messages, where each of these messages also has an expandable node that encapsulates an associated message stack. This technique pushes all the information of importance in this scenario to top-level for quick and efficient analysis. Moreover, because the response message is at top-level, you don't have to search any further to begin analyzing the data. Note that the Operations feature aligns with the overall Message Analyzer design strategy to bring relevant data that is normally hidden or dispersed to the top-level of a suitable analysis surface, where you can examine it immediately without the encumbrance of having to apply elaborate techniques to get at the data you need to assess.

> **NOTE**
> Some typical protocols that support request/response message pairs include HTTP, DNS, and SMB2. The messages for these protocols and others that support Operations will appear in the **Analysis Grid** viewer as Operation nodes, where they are signified by a blue-cubed icon.

Advantages of Operations to Analysis

The major advantage of having request/response pairs encapsulated into separate Operations nodes is that this configuration provides quick access to both the request and response messages when you expand an Operation node. You can also expose the server response time to requests that Message Analyzer automatically calculates for you, as described in the list that follows. With this information exposed at top-level in the **Analysis Grid** viewer, you can quickly make comparisons between several important values that can be critical to troubleshooting and analysis. These values consist of the following:

- **ResponseTime** — provides an indication of how long it took to receive a server's first response to a request message. Large values for **ResponseTime** can be an indication of a slow server that is having performance problems.

  You can view this value by adding **ResponseTime** as a new column in the **Analysis Grid** viewer from the **Global Annotations** node of the **Field Chooser Tool Window**. Under this node, locate the **ResponseTime** annotation, right-click it, and then select the **Add As Column** command in the context menu that appears. Thereafter, you can see at-a-glance the server **ResponseTime** that is associated with the request and response messages that are encapsulated by any Operation node in the **Analysis Grid** viewer.

- **TimeElapsed** — provides an indication of total elapsed time for an Operation that includes how long it took to receive a server's first response to a request message *plus* the amount of time it took to receive all

fragments associated with the transaction. Large values for **TimeElapsed** can be an indication of network latency issues, particularly if the **ResponseTime** values are comparatively low.

You can view **TimeElapsed** values in the **Analysis Grid**, as it is a column in the default layout for this viewer.

> **TIP**
> If you are working with HTTP Operations, you can open a **Chart** with the **Average Response Times for Operations** view **Layout** against a set of trace results to review a summary of average **Response Time** values for HTTP methods. To open a **Chart** with the specified **Layout**, click the **New Viewer** drop-down list on the global Message Analyzer toolbar, highlight **Chart**, and then select the **Average Response Times for Operations** layout.

More Information
To learn more about the Response Time, as used in the **Average Response Times for Operations** view **Layout**, see the Average Response Time for Operations topic.

Disabling Operations

When you capture messages that are part of an Operation, the Message Analyzer parsing process normally identifies and collapses this traffic to combine related request and response message pairs into a single, top-level message row as previously indicated. To indicate an Operation in the **Analysis Grid** viewer, Message Analyzer designates this row with a blue-cubed icon in the **MessageNumber** column and provides an expansion node next to this icon that opens the request and response message pair that makes up the Operation. However, for viewing preferences and because on occasion it might be advantageous to display the original chronological context in which response messages occurred, Message Analyzer enables you to create this display configuration by selecting the **Disable Operations** item from the **Viewpoints** drop-down list on the Filtering Toolbar. To return to the default message display in the **Analysis Grid** viewer with Operation nodes intact, you can simply select the **No Viewpoint** item from this same list.

The default configuration for message lines in the **Analysis Grid** viewer is to show all top-level message nodes in a set of trace results, which includes Operation nodes that encapsulate request/response pairs for protocols with this type of architecture, along with all other top-level message nodes that are not Operations, both of which have nested expansion controls to expose the underlying encapsulated message stacks. This display configuration is shown in the following figure.

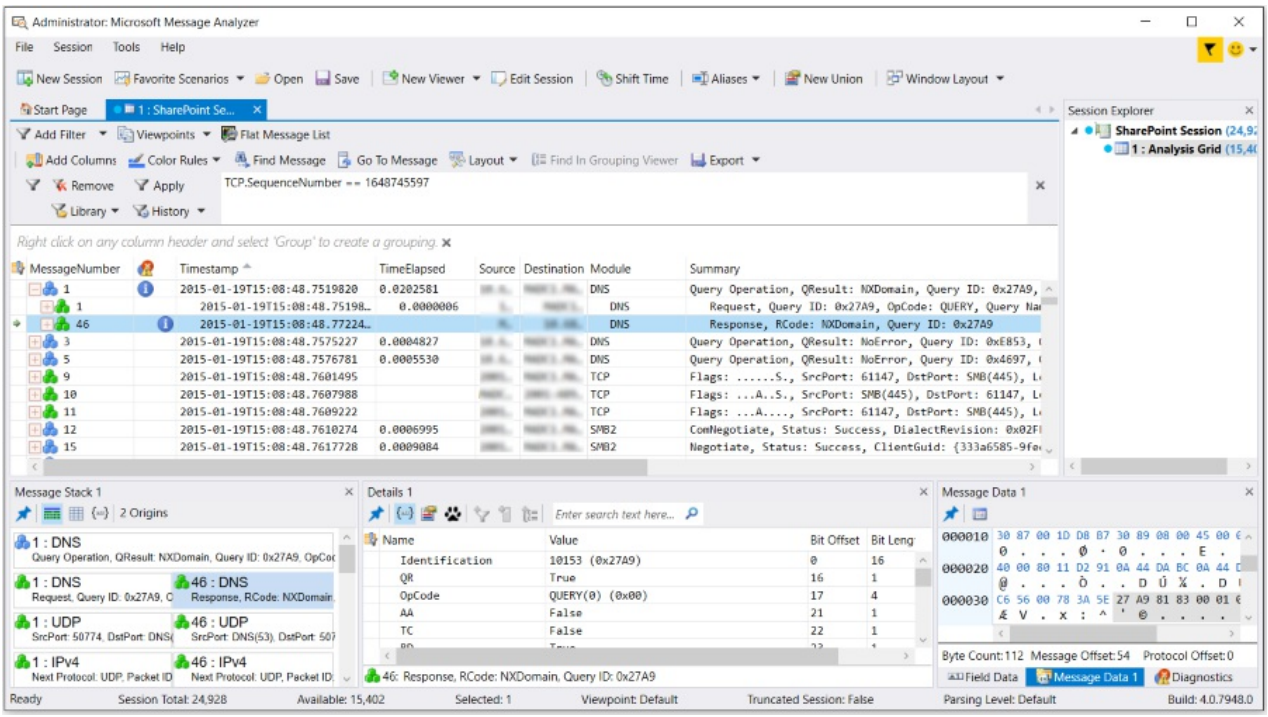

Figure 55: Operations enabled and showing expanded node exposing DNS request/response messages

The figure shows the default view whenever you have **Disable Operations** *unselected* in the **Viewpoints** drop-down list. For example, the **DNS** request and response messages #1 and #46, respectively, are encapsulated in an Operation node that is denoted by a blue-cubed icon, and is currently shown in the expanded state. When you *select* **Disable Operations** in this list, Operation messages are released from their default encapsulated configuration and the constituent request and response messages are broken out and reorganized into top-level messages in their original chronological capture sequence along with all other top-level, non-operation messages in your trace results, as shown in the figure that follows. Note that all top-level messages of the non-operation type are designated by a green-cubed icon.

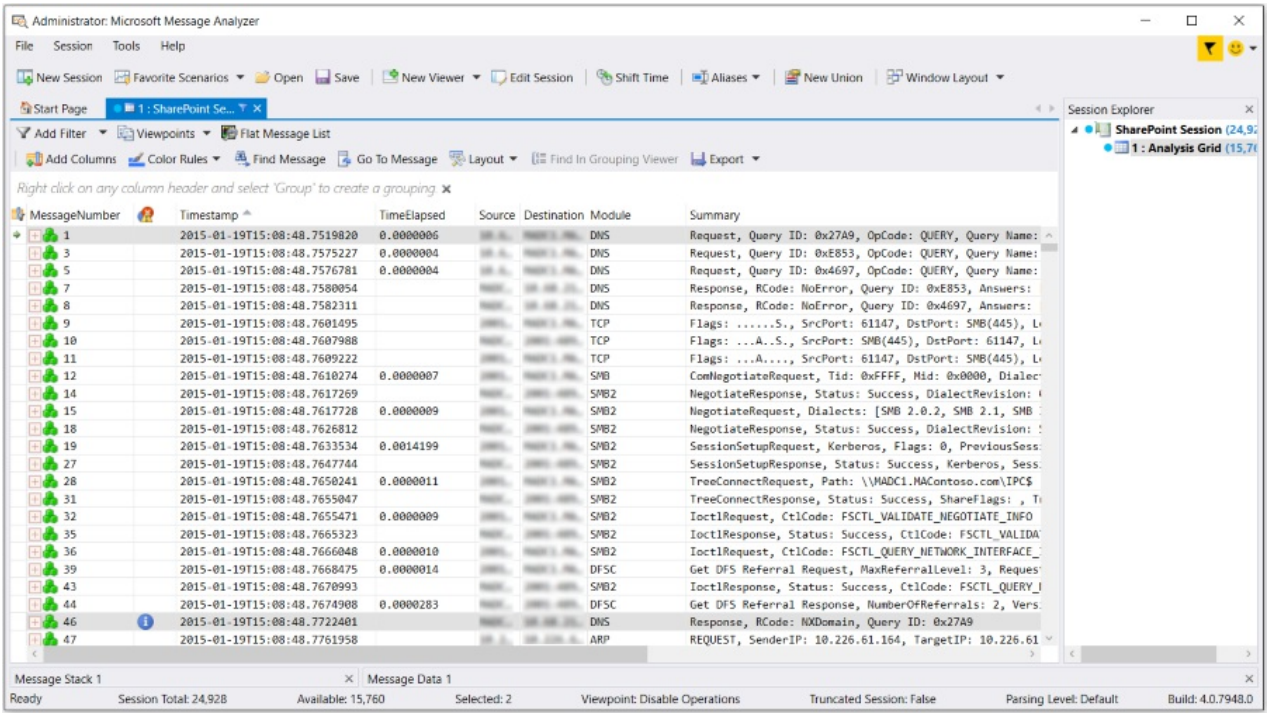

Figure 56: Operations disabled and showing original message sequence

In this figure, note that the formerly encapsulated request message #1 and the response message #46 now display in their original capture sequence.

> **NOTE**
>
> When you **Disable Operations**, Message Analyzer reparses all messages in the Message Store (a repository for the original set of captured but unparsed messages) without encapsulating the request and response message pairs under Operation nodes.

Operations and Filtering Interactions

A situation where you might need to disable Operations is when you want to apply a view **Filter** that isolates a request message type or a response message type from a particular protocol that uses this architecture, such as SMB2. When request and response message pairs are encapsulated under an Operation node, a view **Filter** will be unable to separate them so that you can view them as standalone messages of either type. For example, if you created and applied a **Filter** such as `SMB2.NegotiateRequest`, you will return the SMB2 request messages within the context of the Operation nodes only, rather than as standalone messages as you might typically expect.

The only way that you can isolate these messages is to apply a **Viewpoint Filter** after you have broken the request and response message pairs apart with the **Disable Operations Viewpoint**. Thereafter, you should be able to see the isolated SMB2 request messages, as described in the steps below:

1. Through a Data Retrieval Session, import a message collection that includes SMB2 messages.

   The messages should have been captured while performing some SMB2 file share access operations.

2. Ensure that **No Viewpoint** is currently selected in the **Viewpoints** drop-down list.

3. Type the filter code `SMB2.NegotiateRequest` in the text box of the **Filter** panel on the Filtering Toolbar.

4. Click the **Apply** button on the Filtering Toolbar and note that the results display SMB2 Operation nodes rather than separate request messages, as shown in the figure that follows.

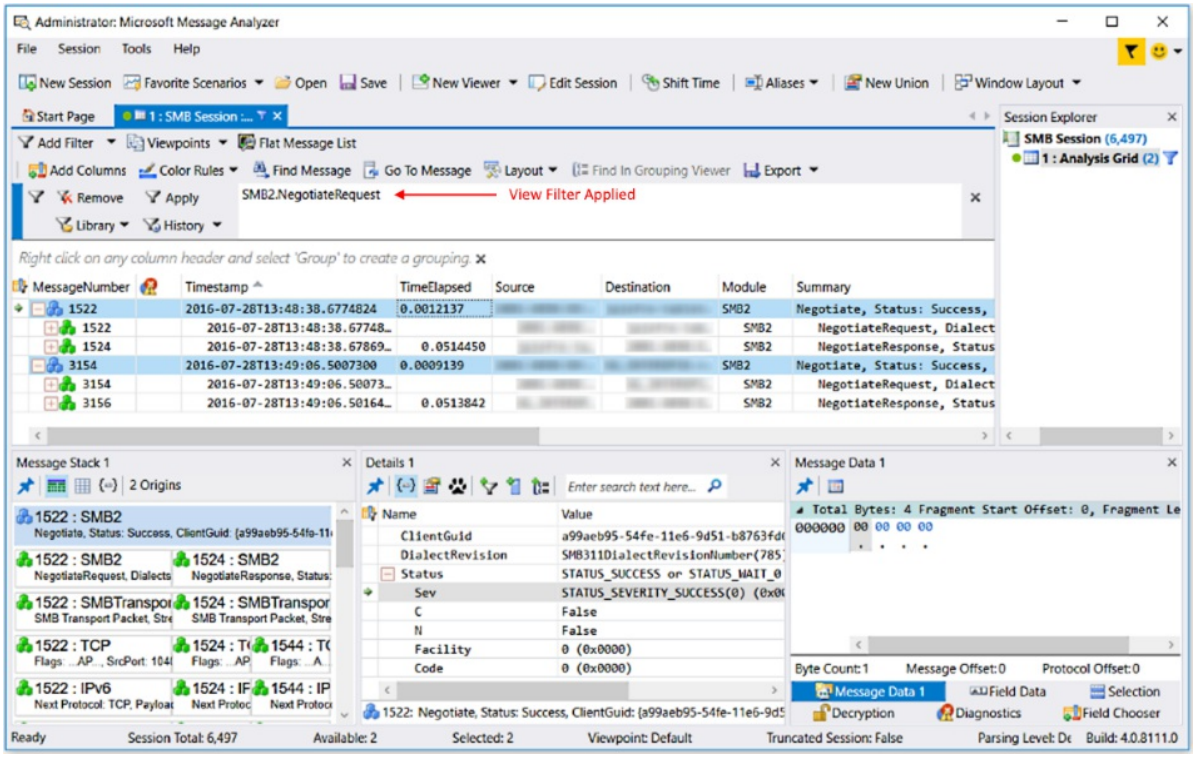

Figure 57: View Filter fails to isolate SMB2 request messages

5. Select the **Disable Operations** item in the **Viewpoints** drop-down list to disable all Operations.

   Note that both SMB2 request *and* response messages continue to display, given that the effect of the view **Filter** is still in play here, as shown in the figure that follows.

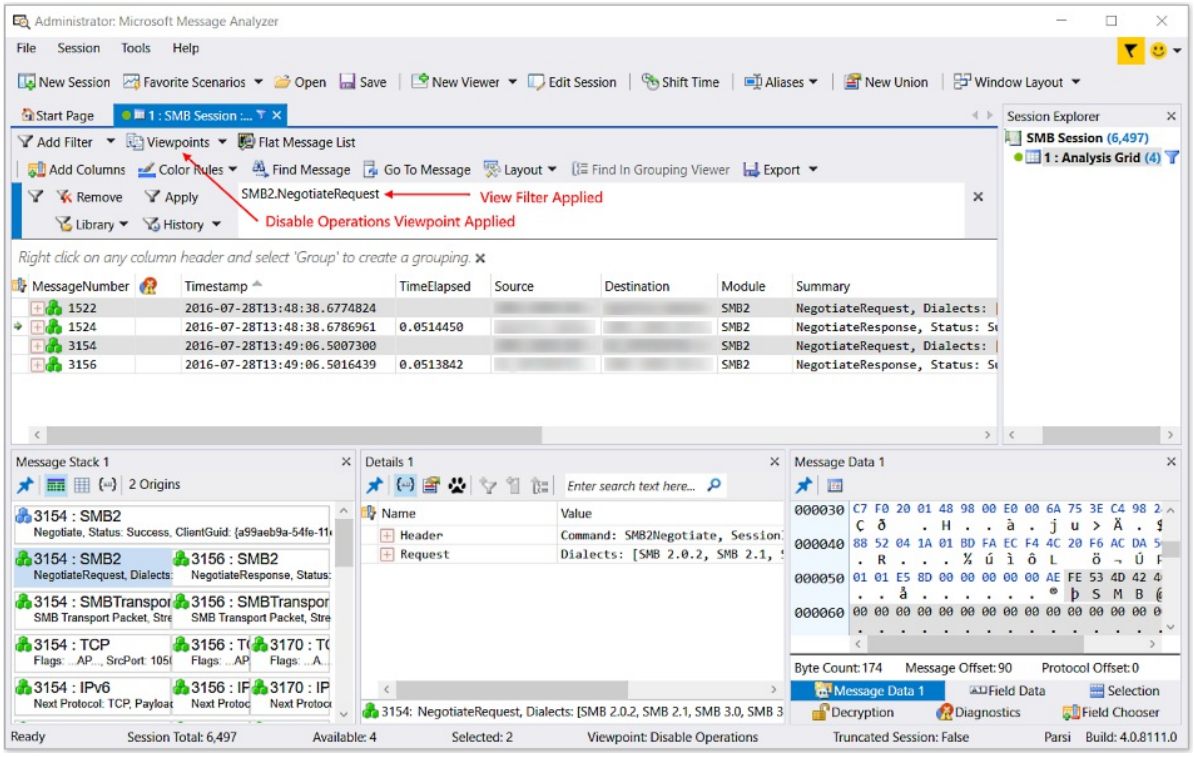

Figure 58: Disable Operations fails to isolate SMB2 request messages

6. Open a **Viewpoint Filter** panel by clicking the **Add Viewpoint Filter** command in the **Add Filter** drop-down list on the Filtering Toolbar.

7. Enter the `SMB2.NegotiateRequest` filter code in the Filter Expression text box of the **Viewpoint Filter** panel and then click the **Apply** button on the panel.

   Message Analyzer now shows only the SMB2 *request* messages in the **Analysis Grid** viewer, as shown in the following figure.

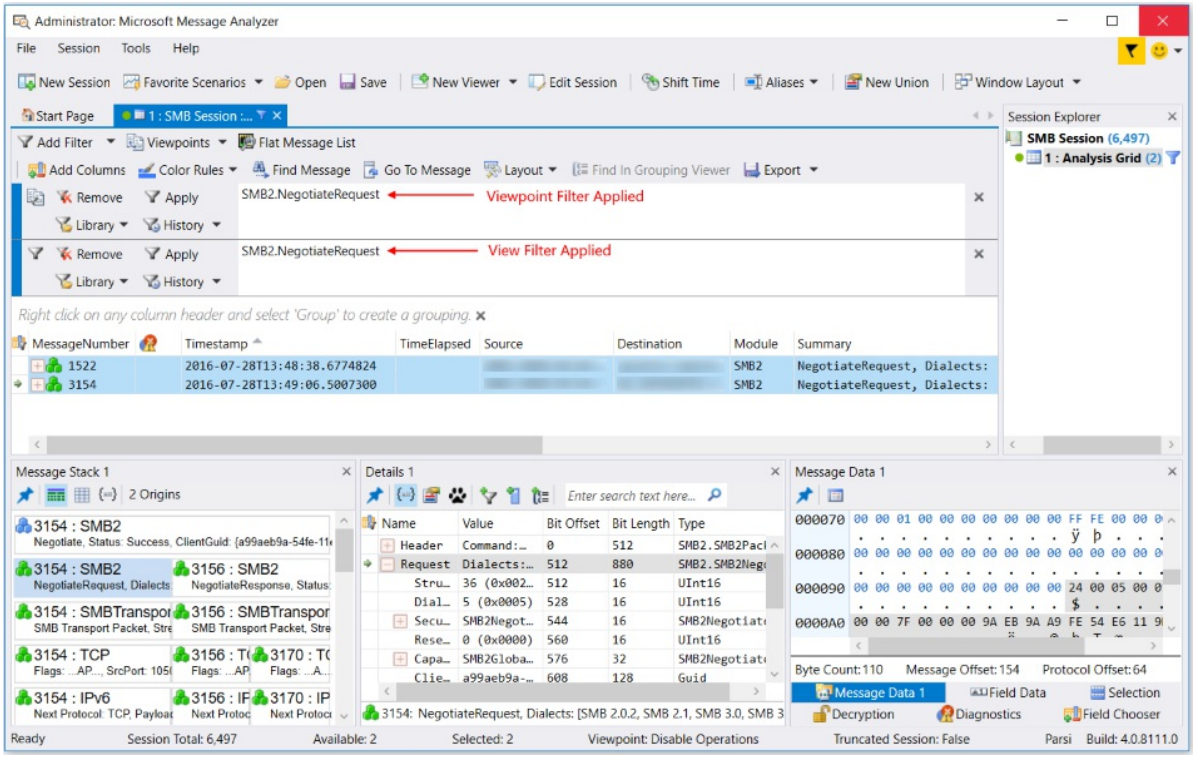

Figure 59: Viewpoint filter succeeds in isolating SMB2 request messages

As a rule of thumb, keep in mind that you can successfully filter under an applied **Viewpoint** only by applying a **Viewpoint Filter**, as a view **Filter** will not provide the desired results in this scenario.

TIP

In the previous procedure, you can substitute selection of the **SMB/SMB2 Disable Operations** command in the **Viewpoints** drop-down list for step 4.

Usage Scenarios with Viewpoints

In a typical usage scenario, you might select the **Disable Operations Viewpoint** in the **Viewpoints** drop-down list to cause Operations to be broken apart and the constituent request and response message pairs to then be displayed in chronological order, as previously described. In this display configuration, you might lose some context as the request and response messages will no longer be grouped together as a single operation, but will instead assume their original chronological position in the trace before Message Analyzer created the Operation nodes. This produces a view that is similar to Network Monitor and may provide some analytical value and familiarity to Network Monitor users, but the response messages can still be difficult to locate. However, you might find it easier to correlate the messages if you select the **No Viewpoints** item in the **Viewpoints** drop-down list to return to the display configuration that collapses all related request and response message pairs back into separate, top-level Operation nodes in the **Analysis Grid** viewer.

As an example of another usage scenario, you may have applied a viewpoint such as the **HTTP Viewpoint** during an Analysis Session because you want to view only the messages at that level. However, HTTP can also have non-operational messages that display at the applied **Viewpoint**, for example payload reassembly messages, and you might not want to examine these yet. To remove these non-operational messages, you can apply the view **Filter** `HTTP && *IsOperation` from the Filtering Toolbar. If you want to examine only the non-operational HTTP messages, apply the view **Filter** `HTTP && !*IsOperation` from another **Filter** panel on the Filtering Toolbar. You can then toggle back and forth between applying and removing these filters to view and analyze the Operation messages and non-Operation messages in the **Analysis Grid** viewer, as required.

TIP

After you select the **Disable Operations Viewpoint** to disassociate the request/response pairs in the Operation nodes for a set of trace results, you can continue to view the request and response message stacks side-by-side in the **Message Stack Tool Window**.

See Also

Using the Filtering Toolbar
Applying and Managing Viewpoints

Creating a Flat Message List

2 minutes to read

A new feature is now included in Message Analyzer for Microsoft Network Monitor users who are making the transition to the Message Analyzer tool, to enable them to create a flattened message display that appears similar to the way Network Monitor displays a set of trace results. Users can initiate this type of message display by clicking the **Flat Message List** button on the Filtering toolbar.

When you select this command, messages are displayed at top-level in their original chronological order, where each request/response message pair is no longer encapsulated under a separate Operation node. However, note that there is a slight difference between Message Analyzer and Network Monitor in this context, in that Message Analyzer's top-level messages also encapsulate the origins (stack) messages, including fragments under expandable nodes, whereas Network Monitor does not. The included message stack data provides an additional level of detail that is quickly accessible for analysis. Moreover, even while the **Flat Message List** command is active against a set of trace results and Operations are broken apart, you can still see the request and response message pairs side-by-side in the **Message Stack Tool Window**, by selecting any request message in the **Analysis Grid** viewer.

After you apply the **Flat Message List** command to a set of trace results and you want to return to the default message display configuration in the **Analysis Grid** viewer with Operation nodes intact, you can simply click the **Flat Message List** button again. Note that while the **Flat Message List** command is active against a set of trace results, its associated button on the Filtering toolbar remains highlighted.

> **NOTE**
>
> When you apply this command to a set of trace results, you may notice a significant difference in message count that displays in the **Available** label on the Message Analyzer status bar. This can result from message fragments that were formerly encapsulated under an Operation node, but are now pushed to top-level in the **Analysis Grid** viewer, thus adding to the **Available** message count.

Filtering Data Sources

4 minutes to read

This section describes how to filter a collection of messages based on the associated data sources from which the messages derive. This feature is particularly useful if you are loading saved data from multiple sources into Message Analyzer through a Data Retrieval Session. For example, you might have a collection of messages from related sources such as logs and traces that you are correlating and it may be advantageous to isolate the messages that are associated with each source of data. There are two ways you can achieve this, as follows:

- **Filtering** — filter on the data sources that contain the messages you want to view.

  To do this, select one or more data sources in the **Data Source Filter** dialog that is accessible from the global Message Analyzer **Session** menu after you load a collection of messages into Message Analyzer. By selecting one or more data sources in the dialog and clicking the **Apply** button, all messages are filtered from the current data viewer except those that are associated with the selected data source/s.

> **NOTE**
>
> The **Data Source Filter** feature does not support filtering in **Layouts** for the **Chart** viewer that you can access from the **New Viewer** drop-down list.

The **Data Source Filter** dialog is shown in the figure that follows, where a trace file and a cluster log are the selected data sources.

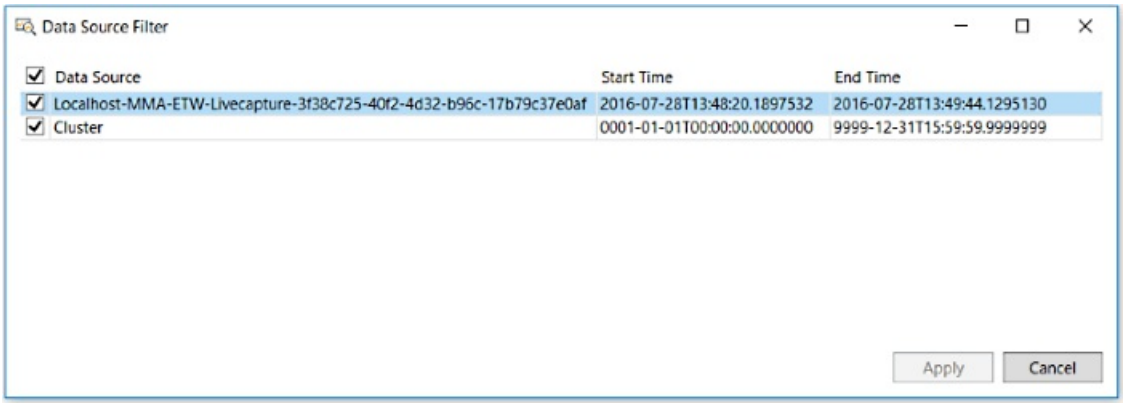

Figure 60: Data Source Filter dialog

- **Grouping** — create an associated message Group for each data source in a message collection under separate expandable **Data Source** nodes in the **Analysis Grid** viewer.

  To do this, add a **DataSource** column to the **Analysis Grid** viewer by right-clicking **DataSource** under the **General** node of the **Field Chooser Tool Window**, and then select the **Add as Column** command in the context menu that appears. Next, right-click the newly added **DataSource** column in the **Analysis Grid** viewer and select the **Group** command in the context menu that appears. The messages from each data source are then organized and encapsulated under a separate node that you can expand for analysis, as shown in the figure that follows.

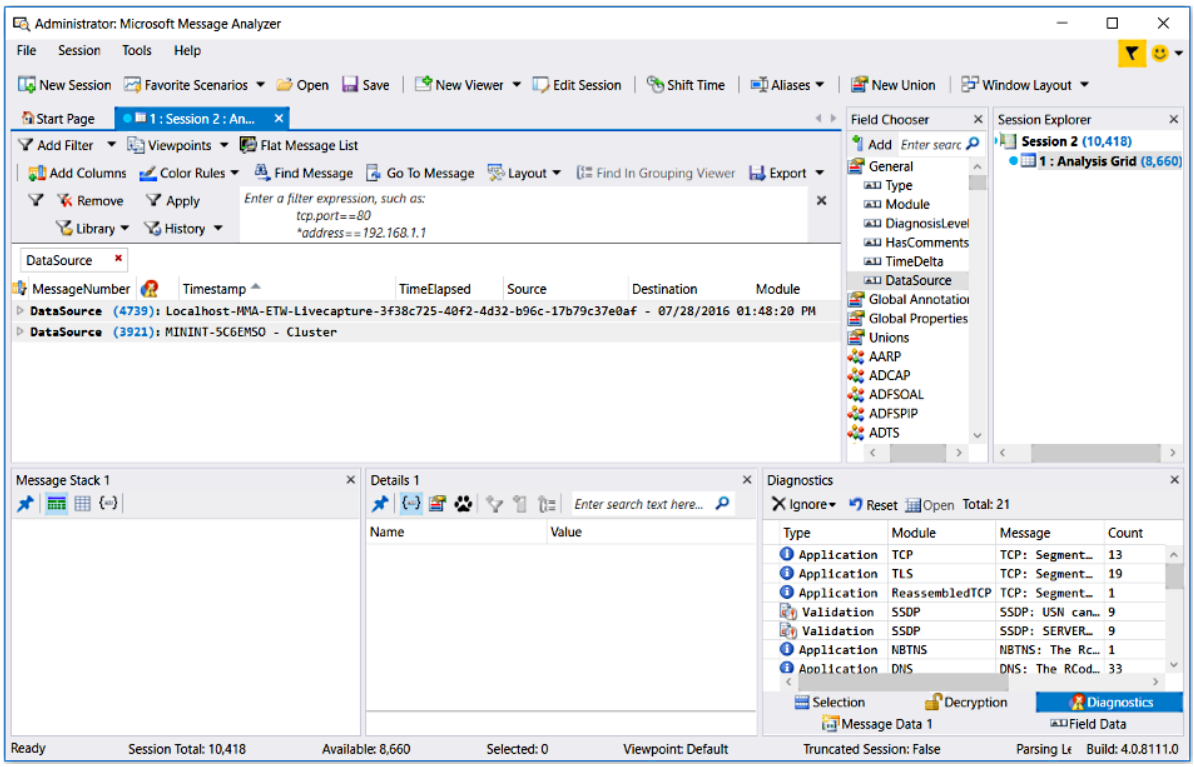

Figure 61: Data Source Groups in the Analysis Grid

> **TIP**
>
> The **Grouping** viewer does not currently have any **Layouts** that make use of the **DataSource** field by default. However, you can manually add this field to any **Grouping** viewer **Layout** as needed, as long as the **Grouping** viewer *has focus* when you are doing so. To add the **DataSource** field to a **Grouping** viewer **Layout**, use the **Field Chooser** as described immediately above.

Using the Data Source Filter Dialog

As previously described, you can access the **Data Source Filter** dialog by clicking the global **Session** menu, highlighting the **Data Source Filter** item, and then selecting the **Edit** command that appears in the drop-down list. When the **Data Source Filter** dialog displays, you will see a tabular listing of data source information in the following column headers:

- **Data Source** — the listings in this column are specified in a format similar to the following examples, for the indicated file types:
  - Message Analyzer Trace Parsed (.matp) files — uses the format: HostName-MMA-ETW-LiveCapture-GUID.
  - Capture (.cap) files — uses the format: FileName
  - Event Trace Log (.etl) files — uses the format: FileName
  - Log (.log) files — uses the format: FileName

> **NOTE**
>
> The **DataSource** column in the **Analysis Grid** viewer can also include additional information when messages from cap, .etl, and other files are displayed. For example the host name can be included in the listing.

- **Start Time** — specifies the **Timestamp** of the first message in the trace file or log.

- **End Time** — specifies the **Timestamp** of the last message in the trace file or log.

Applying and Removing Data Source Filtering

After you use the **Data Source Filter** dialog to select the data source/s containing the messages that you want to view, click the **Apply** button in the dialog to initiate the filtering process. After you apply the filtering, the **Apply** command is disabled and the **Remove** command is enabled in the **Data Source Filter** drop-down list so you can remove the applied filter as necessary. After you click the **Remove** command, the **Apply** command then re-enables in the **Data Source Filter** drop-down list so that you can re-apply the current data source filtering configuration. You can toggle back and forth between applying and removing the **Data Source Filter** as many times as you want without changing the current data source filter configuration. To change the configuration, simply select the **Edit** command in the drop-down and reconfigure your **Data Source Filter** as required before you re-apply it.

More Information

To learn more about how to use the **Group** context menu command in the **Analysis Grid** viewer, see Using the Analysis Grid Group Feature.

To learn more about the **Grouping** viewer, see the Grouping Viewer topic.

Setting Time Shifts

8 minutes to read

Message Analyzer provides a **Shift Time** dialog that enables you to change the timestamp of captured messages displaying in an Analysis Session. The purpose of this feature is to compensate for skewed system clock values or time zone differences across different computers when comparing traces and logs from those computers, either side-by-side in separate Live Trace Sessions or when viewing interlaced message configurations in a single Data Retrieval Session. Ensuring that traces are chronologically aligned is important for troubleshooting. For example, if you have different traces that are interlaced in a Data Retrieval Session, messages that are offset by skewed system clock values will appear in an incorrect sequence that makes data comparison very difficult.

The following are two different ways that you can apply a **Time Shift** to a set of data sources in a session, along with the circumstances in which you might apply them:

- **Explicitly enter a Time Shift** — in this case, you might not have looked at the content of your data sources, but you know in advance that the **Timestamp** values are off, for example, because some of the data was captured in a different time zone. An explicit time shift value is applied to all messages in selected data sources for the current session.
- **Adjust the TimeStamp of a particular message** — this automatically initiates a **Time Shift** calculation for all messages in selected data sources for the current session. In this case, you are in analysis mode and you discover that message **Timestamps** are off and you need to synchronize certain data sources.

Applying a Time Shift

Before you can apply a time shift, you must open the **Shift Time** dialog in either of the following ways:

- By clicking the **Shift Time** button on the global Message Analyzer toolbar.
- By clicking the **Shift Time** item in the **Shift Time** submenu of the global Message Analyzer **Session** menu.
- Right-clicking any message in the **Analysis Grid** viewer and selecting the **Shift Time...** item in the context menu that appears.

After you open the **Shift Time** dialog shown in the figure below, you can apply a time shift that changes the **Time Stamp** values for a set of messages by configuring a time shift with the **Shift Time** dialog controls and then clicking the **Apply** button. In the figure below, a 3-hour time shift is being added to a data source to account for a time zone difference.

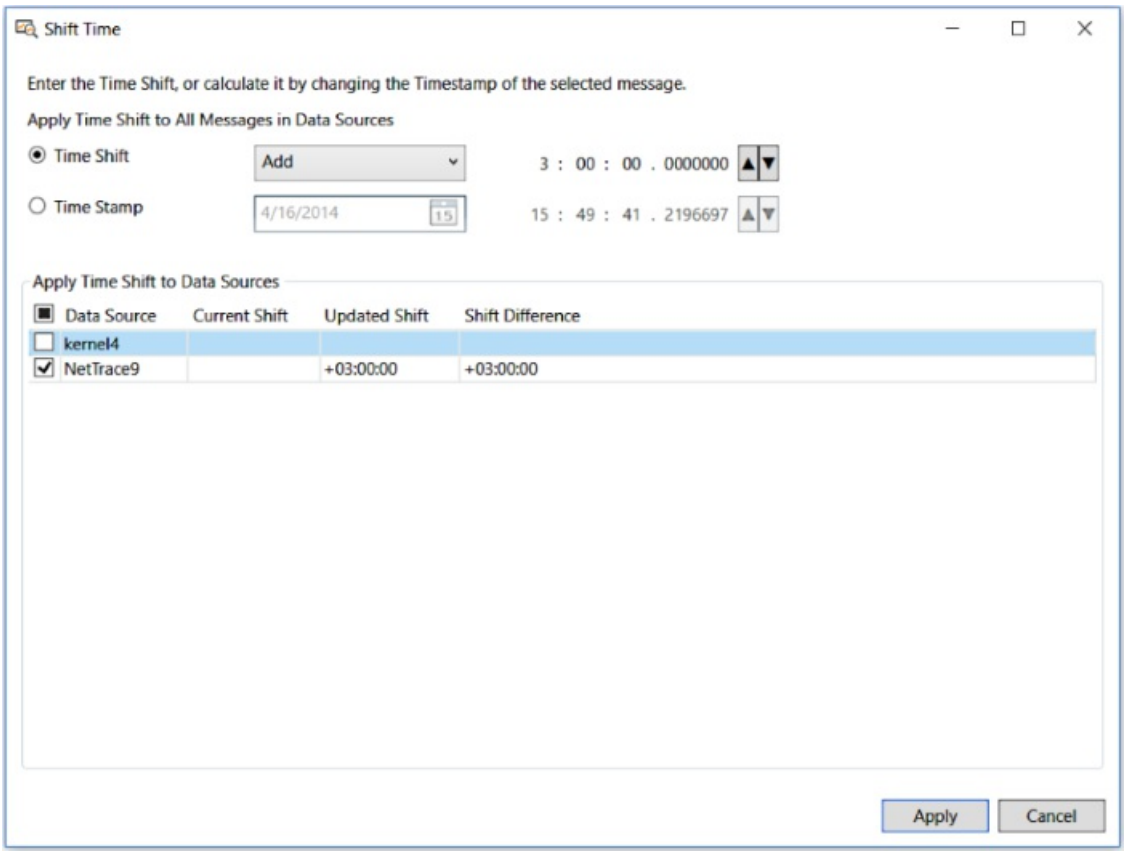

Figure 62: Message Analyzer Shift Time dialog

Dialog Entry Contexts

There are two different contexts in which you can open the **Shift Time** dialog, as follows:

- **Without message context** — if you open the **Shift Time** dialog with no messages selected in the **Analysis Grid** viewer, the dialog is entered *without message context*, where the **Time Stamp** option and related controls are disabled and contain no prepopulated values. This context automatically selects the **Time Shift** option when the **Shift Time** dialog opens, which provides the controls to configure an explicit time shift.

  You can set an explicit time shift to compensate for **Timestamp** differences, by incrementing or decrementing all message **Timestamps** for selected data sources, to adjust for hour, minute, and second value displacements with up to 7 decimal digits resolution. You might use this context because you know that the messages from a particular data source need a time shift and you know what that value is.

- **With message context** — if you open the dialog while a message is selected in the **Analysis Grid** viewer, the dialog is entered *with message context*, where the **Time Stamp** option controls are enabled and contain prepopulated values associated with the selected message. This context provides options to configure a time shift based on the **Timestamp** value of a selected message, where you can accommodate for message **Timestamp** differences between data sources by specifying settings that calculate the following:

  - A change to the date of all message **Timestamps** in selected data sources.
  - An incremental or decremental **Timestamp** change to all messages in selected data sources, by making adjustments to hour, minute, and second value displacements with up to 7 decimal digits resolution.

You might be using this context because you discovered during analysis that messages from a particular data source need a time shift and you want to calculate the shift based on a known/selected message **Timestamp** value.

Context-Enabled Dialog Controls

When you open the **Shift Time** dialog *without message context* by clicking the **Shift Time** drop-down list on the global Message Analyzer **Session** menu, the following menu items are available to perform the indicated actions (these options are unavailable when you click the **Shift Time** button on the global Message Analyzer toolbar):

- **Shift Time** — displays the **Shift Time** dialog.

  When you open the dialog without message context, the **Time Shift** option and associated controls are enabled, with the **Time Shift** drop-down set to **Add** and the **Time Shift** spin control values set to zero. The **Time Stamp** option is disabled in this context, but the **Apply Time Shift to Data Sources** pane may be populated with values if you previously specified a time shift in the current session. By using the **Time Shift** controls that display in the *without message context* mode, you can specify an explicit incremental time shift for all messages in selected data sources to compensate for skewed **Timestamp** values.

- **Remove All Time Shifts** — removes all time shifts for the current in-focus session.

When you open the **Shift Time** dialog *with message context*, that is, while a message is selected in the **Analysis Grid** viewer, all controls in the **Shift Time** dialog are enabled so you can specify an incremental time shift value as well as a date change to accommodate **Timestamp** differences. Also, the **Apply Time Shift to Data Sources** pane may be populated with values if you previously specified a time shift and reopened the **Shift Time** dialog in the current session.

Using the Time Shift Controls

The controls in the **Shift Time** dialog provide time shift functionality as follows:

- **Time Shift** option — enables the following controls that allow you to set an incremental time shift value — in hours, minutes, and seconds, with up to 7 decimal digits resolution — that alters all message **Timestamps** in selected data sources:
  - **Time Shift** drop-down — sets the arithmetic operator for the time shift, by specifying either the **Add** or **Subtract** menu item.
  - **Time Shift** up-down — sets the actual incremental time shift value that you apply to selected data sources.
  - **Time Shift** spin — works interactively with the **Time Shift** up-down control by enabling you to place your mouse cursor in any of the spin control fields, consisting of hh:mm:ss.ddddddd, and click the up or down arrows to add or subtract a specified increment of time, respectively.
- **Time Stamp** option — enables the following controls to allow you to set an incremental time shift value and/or change the calendar date for all messages in selected data sources:

> **NOTE**
> If you enter the **Shift Time** dialog *with message context*, the **Time Stamp** controls reflect the date-time stamp values of the message that is currently selected in the **Analysis Grid** viewer.

- **Date** — consists of a date control that drops down when you click it. Enables you to specify a date that shifts the **Timestamp** values of all messages in selected data sources in 24-hour increments.
- **Time Stamp** up-down and spin — works interactively to enable you to add or subtract a specified increment of time, as previously described.
- **Apply Time Shift to Data Sources** pane — specifies time shift statistics that include **Current Shift**,

Updated Shift, and **Shift Difference** values. Also specifies the data sources for the current session in the **Data Source** column.

Time Shift Example

In practice, it is likely that you will apply a **Time Shift** to one or more selected data sources. For example, you might have loaded data from two trace files and you either already know that a time shift is required for one of the files, or you discover this during analysis.

In the first case, where you did not select any messages in the **Analysis Grid** viewer (without message context), you can open the **Shift Time** dialog in the previously stated manner and simply select the **Data Source** in the **Apply Time Shift to Data Sources** grid of the dialog for the messages that require a shift, while deselecting all others. You can then use the **Time Shift** controls to add or subtract a configured time shift value and click the **Apply** button to increment or decrement the **Timestamps** for the selected data source by the value that you specified. This action applies the time shift to all messages in a selected data source. At this point, you can sort the **Timestamp** column in the **Analysis Grid** viewer to interlace the message collection in chronological order for analysis purposes.

In the second case, where you select and right-click a particular message in the **Analysis Grid** viewer (with message context) and select the **Shift Time** menu item, the **Shift Time** dialog opens with the **Time Stamp** option/controls enabled and containing the **Timestamp** value of the selected message. You can specify whether to add or subtract a **Time Shift** value in the **Time Shift** drop-down and you can now select the **Time Stamp** option to enable the **Time Stamp** date, spin, and up-down controls. Using these controls, you can specify a date and incremental time shift value based on the value of the selected message. For example, you might want to match the **Timestamp** value of one message to another message from a different data source. Next, select the **Data Source** to which the shift will be applied and deselect the source containing messages that you do not want to shift. When you click the **Apply** button, the time shift value that you specified based on the selected message will be applied to all messages in the selected **Data Source**. Again, you should sort the **Timestamp** column in the **Analysis Grid** viewer at this point, to interlace the message collection in chronological order for analysis purposes.

Configuring Time Format Settings

2 minutes to read

For the data viewers that display date-time values, Message Analyzer utilizes the standard ISO format for consistency throughout the user interface (UI). For example, the same date-time format that displays in the **Analysis Grid** viewer will also display in the **Pattern Match** viewer, in addition to the **Time Filter** dialog in Live Trace Session results and in the **Time Filter** pane of a Data Retrieval Session. However, Message Analyzer enables you to change how time data displays by providing you with the option to show either the date and time, or the time only. Configuration settings for these formats are available from the following locations:

- **Options dialog** — this dialog is accessible from the global Message Analyzer **Tools** menu. The dialog exposes the **Time Display** pane on the **Display** tab, from where you can specify date and time configuration settings that include **Show Date And Time** and **Show Time Only** options, along with a **Time Zone** drop-down list that enables you to set your locale. Specifying the date and time configuration settings from this location has global scope that overrides any previously set configuration. Note that by default, your local time zone is automatically set in the **Time Zone** drop-down list when you start Message Analyzer and will remain that way until you manually change it.
- **Analysis Grid Timestamp column context menu** — right-click the **Timestamp** column header in the **Analysis Grid** viewer to display a context menu that enables you to toggle the date and time configuration settings between **Show Time Only** and **Show Date and Time**; note that the latter is the default. Only one of these commands will display at any time, depending on the current configuration setting. Setting one of these commands from the **Timestamp** column context menu also has global scope and overrides any previously set configuration.

You can set the time locale as the reference time zone for all Data Retrieval and Live Trace Sessions, to achieve a time zone perspective that enables you to view time information based on where your data was captured. This makes it convenient to analyze data based on the perspective of a particular time zone. Also, by setting the time configuration to **Show Time Only**, you can condense the time value display and recover some UI display space in the **Timestamp** column and other viewers.

> **NOTE**
>
> Although you might set the time display configuration to **Show Time Only**, you will still need to specify a full date and time value when you configure a `#Timestamp` filter, for example, when specifying a **Session Filter** or view **Filter**.

See Also

Setting Message Analyzer Global Options

Using and Managing Message Analyzer Aliases

2 minutes to read

Message Analyzer provides an **Aliases** function that enables you to substitute friendly names for several types of data field values in the **Analysis Grid** and other viewers. **Aliases** facilitate easy recognition of values that otherwise can be somewhat cryptic and difficult to work with. Message Analyzer maintains a default set of **Aliases** and any new ones that you create in an **Aliases** Library, which is accessible from the **Aliases** drop-down list on the global Message Analyzer toolbar and from the global Message Analyzer **Tools** menu. The intent of this feature is to improve your ability to discover and analyze specific message traffic through the use of simple user-defined naming conventions that have meaning in the context of your troubleshooting environment. By customizing your data analysis environment with aliasing, you can make it easier to keep track of traffic to or from different host IP addresses, physical addresses, ports, and so on. The data field types that currently support aliasing consist of the following:

- **Source and Destination IP Address** — for any **Source** or **Destination** IP address that displays in the **Analysis Grid** viewer, you can substitute a friendly name for the address value, for example, "MyComputer".
- **Source and Destination Media Access Control (MAC) Address** — for any message that can display a MAC address data field value in the **Analysis Grid** viewer, for example an ARP message, you can substitute a friendly name for the address value, for example, "SenderAdapterMAC", "TargetAdapterMAC", and so on.
- **TCP SourcePort and TCP DestinationPort** — for any message that can display a TCP **SourcePort** or **DestinationPort** data field value in the **Analysis Grid** viewer, you can substitute a friendly name for the port. For example, you might use an application acronym prefix in the alias name that corresponds with the associated port, such as "HTTPPort", "DNSPort", "LDAPPort", and so on.

> **NOTE**
>
> In the case of MAC addresses, TCP **SourcePort**, and TCP **DestinationPort**, it is likely that you will need to use the **Field Chooser Tool Window** to add the necessary columns in the **Analysis Grid** viewer for displaying the values of these data field types, before you can create an **Alias** for a value in one or more of these columns.

Alias Example

As an example, the IPv6 address of a server, such as `FE80:0:0:0:4D45:3FCD:BDE0:69BE`, can be very difficult to read and keep track of when performing data analysis, as it is difficult to distinguish from other IPv6 addresses. Users often have to create an alternate method of externally mapping these addresses to a more friendly name, for example in a Notepad listing, which can be time-consuming and cumbersome during data analysis. With the Message Analyzer **Aliases** feature, you can simply replace such an IPv6 value with a friendly name such as "WebServer" or "DatabaseServer" to make traffic from the associated **Alias** address immediately obvious.

What You Will Learn

In the following topics of this section, you will learn how to create, manage, share, and perform Message Analyzer operations with aliases:

Creating Message Analyzer Aliases

Modifying Message Analyzer Aliases

Enabling and Disabling Message Analyzer Aliases

Performing Message Analyzer Operations with Aliases

Creating Message Analyzer Aliases

5 minutes to read

Message Analyzer enables you to create an alphanumeric string value **Alias** for any data field value in the **Analysis Grid** viewer that is of a type that supports aliasing, as described in Using and Managing Message Analyzer Aliases. You can create an **Alias** for only one data field value at a time; however, you can only create an **Alias** from the **Analysis Grid** viewer, as there are no other facilities to create one. Note that you can apply any alias while Message Analyzer is capturing data in a Live Trace Session or loading messages through a Data Retrieval Session.

Go To Procedure
To go directly to a procedure that creates an **Alias**, see Create an Alias for a Data Field Value. However, you are advised to examine the information contained in this section before doing so.

Enforcing Unique Alias Values

Message Analyzer enforces the restriction that the **Value** of each *applied* **Alias** must be unique. However, if you create an **Alias** and its **Value** matches that of another **Alias**, and you **Save** it, Message Analyzer simply adds that **Alias** to the **Aliases** drop-down list on the global Message Analyzer toolbar and to the **Aliases** submenu that is accessible from the global Message Analyzer **Tools** menu, and then sets it to the disabled state, while preventing it from being applied to the current message set. If you attempt to enable this **Alias**, Message Analyzer blocks application of the **Alias** and displays a **Duplicated value** message to indicate the inherent restriction. Even though this is a restriction, it does give you the flexibility to create multiple **Aliases** with different names for the same **Value**; however, you can only enable and apply them one at a time.

> **NOTE**
> Note that although it is possible to create multiple aliases with the *same* name but with different **Values**, this is probably not the best use of this feature, except perhaps for multiple IP addresses that your computer may have.

Configuring an Alias

To configure an **Alias** for a particular data field value in a set of trace results, right-click a field value in the **Analysis Grid** viewer that supports aliasing, for example, an IP address in the **Source** column, and then select the **Create 'Source' Alias** context menu item. Thereafter, the **Alias Editor** dialog appears and enables you to specify the values or settings that follow:

> **NOTE**
> The **Create '<columnName>' Alias** menu item is enabled only for data field values that support aliasing. In addition, after you create a new **Alias**, the **Create '<columnName>' Alias** menu item for the new **Alias** will be grayed-out and disabled (if you right click the new **Alias**) until such time that you delete the existing **Alias** from the **Aliases** drop-down list. Note that the value '*<columnName>*' is the name of the **Analysis Grid** viewer column that contains data that is of a type that currently supports aliasing, for example, the **Destination** column that contains IP addresses.

- **Value** — the value of the data field that you right-clicked in the **Analysis Grid** viewer to create a new **Alias** is automatically passed in to the **Value** text box. For example, this could be an IP address or a TCP port number. This field is editable but cannot be blank, although Message Analyzer performs a validation that

determines whether the entered value is correct for the field type.

- **Alias** — enter a friendly name for the new **Alias** in this text box. Specify a name that makes sense for your environment. This field is editable but cannot be blank; it should be between 1-64 characters long. Note that the alphanumeric string that you specify as the **Alias** name is automatically interpreted as a string in various Message Analyzer operations such as filtering, sorting, and grouping.
- **Description** — type a description of the new **Alias** in the **Description** text box, to define the usage context for future recollection and to notify users with whom you might share such an **Alias** asset, as the **Description** is included in the Message Analyzer sharing infrastructure. You also have the option to leave this text box blank.
- **Category** — in this text box, specify the name of the **Category** in which to store your new **Alias**. You can choose an existing **Category** that is populated in this combo box based on previously specified **Categories**, or you can create your own customized **Category** name, which is then retained as an editable **Category** drop-down list item for future use. Thereafter, you can simply select items from the **Category** drop-down menu as needed. If you create your own **Categories**, you might consider naming them by data types, for example, "IP Addresses", MAC Addresses", "TCP Ports", and so on. All categories that you specify become subcategories that appear under the default **My Items** top-level category in the **Aliases** drop-down list. However, if you leave the **Category** combo box blank, the **Alias** will be stored directly in the **My Items** category.
- **Auto Refresh Views** — select this check box if you want Message Analyzer to refresh all data viewers that will be impacted by the new **Alias**, immediately after you **Save** the **Alias**.

> **IMPORTANT**
>
> Be aware that the **Auto Refresh Views** function of the **Alias Editor** dialog can be activated only when creating a new **Alias**. It does not activate for other operations that you can perform, such as editing and saving an **Alias** or creating a copy of, modifying, and saving an **Alias**. In addition, when the current message set is refreshed to apply a newly created **Alias**, there could be an impact on performance depending on the number of messages that exist in the displayed message set.

If you do not select the **Auto Refresh Views** check box, and you **Save** the new **Alias**, Message Analyzer prompts you with the following message that displays on the **Aliases Changed** information bar above the **Analysis Grid** viewer:

"**To avoid excessive CPU utilization when refreshing some data views, please wait until all alias changes are complete before you refresh.**"

Thereafter, you have the option to manually click the **Refresh Views** button on the information bar at a time when it is appropriate to refresh your data viewers.

When you are finished entering the **Alias** data in the **Alias Editor** dialog, click **Save** to store your new **Alias** in the specified category of the **Aliases** drop-down list. After adding the new **Alias** to your **Alias** collection, you can edit, make a copy, or delete the **Alias** at any time, as described in Modifying Message Analyzer Aliases.

> **NOTE**
>
> Your Message Analyzer installation will save any **Aliases** that you create for use in subsequent sessions, however, **Aliases** do not persist in any data files that you save and share with others, such as a .matp trace file.

Modifying Message Analyzer Aliases

4 minutes to read

If you want to modify an existing **Alias**, you will need to open the **Alias Editor** dialog, which is the same dialog that enables you to create a new **Alias**. To backtrack for a moment, when you *create* a new **Alias**, you access the **Alias Editor** dialog by right-clicking the field value of a data type in the **Analysis Grid** viewer that supports aliasing and then selecting the **Create '<columnName>' Alias** menu item to display the dialog. However, to *modify* an existing **Alias**, you can only use either of the following methods to open the **Alias Editor** dialog:

- Right-click an existing **Alias** in the **Aliases** drop-down list that is accessible from the Message Analyzer global **Tools** menu or the global toolbar, and then select the **Edit** item in the context menu that appears.
- Click the **Aliases** drop-down list and then click the **Manage Aliases** command. From the **Manage Alias** dialog that displays, right-click the existing **Alias** that you want to modify, and then select the **Edit** item in the context menu that appears.

Using the Context Menu Commands

Whenever you right-click an existing **Alias**, either in the **Aliases** drop-down list or in the **Manage Alias** dialog, the following commands are accessible from the context menu that displays:

- **Edit** — click this command to open the **Alias Editor** dialog, as previously indicated. The values and settings that you can specify in the **Alias Editor** dialog are described in Creating Message Analyzer Aliases. Note that after you edit and **Save** changes to a currently enabled **Alias**, your modifications are automatically applied to the current message set, without Message Analyzer displaying the **Aliases Changed** information bar that prompts you to refresh data views. If you edit and **Save** changes to a currently disabled **Alias** (see Enabling and Disabling Message Analyzer Aliases), your modifications are saved but they are not applied to the current message set until you manually select the **Alias** to enable it.
- **Create a Copy** — click this command to open the **Alias Editor**, where you can work with an exact replica of an existing **Alias**, to either **Save** as is or modify as necessary. By providing an **Alias** template, this command enables you to quickly create a new **Alias** based on similar data from an existing one. When you select the **Create a Copy** command and edit the copy of an existing **Alias**, the following guidelines may apply:
  - **Enforcing Unique Alias Values** — Message Analyzer enforces the restriction that the **Value** of each *applied* **Alias** must be unique, as previously specified. However, if you **Create a Copy** of an **Alias** and modify it such that its **Value** matches that of another **Alias**, and you **Save** it, Message Analyzer simply adds that **Alias** to the **Aliases** drop-down list and sets it to the disabled state, while preventing it from being applied to the current message set. If you attempt to enable this **Alias**, Message Analyzer blocks application of the **Alias** and displays a **Duplicated value** message to indicate the inherent restriction.
  - **Saving Exact Alias Copies** — if you **Save** an **Alias** *copy* as-is, it is added to the **Aliases** drop-down list, but it is not enabled, again because Message Analyzer enforces the restriction that the **Value** of each **Alias** must be exclusive.
  - **Saving Edited Alias Copies** — if you modify a copied **Alias** and **Save** it with an **Alias** name that is different than any other **Alias** name, but its **Value** is the same as another existing **Alias**, the saved **Alias** is added to the **Aliases** drop-down list and is disabled. This gives you the flexibility to create multiple **Aliases** with different names for the same **Value**; however, you can only enable and apply them one at a time.

If you modify a copied **Alias** and **Save** it with an **Alias** name that is identical to an existing **Alias** name, but its **Value** is different than that of any other existing **Alias**, Message Analyzer adds it to the **Aliases** drop-down list in the disabled state. However, you can enable and apply this **Alias** by simply selecting it in the drop-down list, at which time Message Analyzer displays the **Aliases Changed** information bar to prompt you to refresh your data views. After you click the **Refresh Views** button on the information bar, the **Alias** is then applied to the current message set, which includes all sessions and viewers.

- **Delete** — click this command to permanently remove an **Alias** from the **Aliases** drop-down list. If you **Delete** an **Alias** that is the only one in a particular **Category**, the **Category** will be removed as well. Whenever you delete a currently enabled **Alias**, Message Analyzer displays the **Aliases Changed** information bar to prompt you to refresh your data views. After you click the **Refresh Views** button on the information bar, the **Alias** is then applied to the current message set, which includes all sessions and viewers.

Enabling and Disabling Message Analyzer Aliases

3 minutes to read

You can enable or disable any **Alias** by toggling (selecting/unselecting) the appropriate check box in the **Aliases** drop-down list in the global Message Analyzer **Tools** menu or on the global toolbar. You can either enable each **Alias** one at a time, or you can enable one or more **Categories** that each contain multiple **Aliases** at once. When you enable one or more **Aliases** by selecting them, the current message set will be updated such that all data field values that are represented by the **Alias** configurations will be replaced by the **Alias** names. When you disable one or more **Aliases** by unselecting them, each **Alias** is removed from all previously set instances in the current message set that is displayed in the **Analysis Grid** and other viewers. However, each disabled **Alias** is still retained in the **Aliases** drop-down list until such time that you delete it.

> **NOTE**
>
> The scope of **Alias** application encompasses all data viewers and sessions displayed in the current running instance of Message Analyzer, which in this context is referred to as the current message set.

Whenever you enable or disable an existing **Alias**, you are prompted by an **Aliases Changed** information bar that contains the following advisory message:

"**To avoid excessive CPU utilization when refreshing some data views, please wait until all alias changes are complete before you refresh.**"

When this information bar displays, you have the option to refresh your data viewers at the appropriate time by clicking the **Refresh Views** button on the bar.

> **NOTE**
>
> If you enable an **Alias** that has a **Value** that matches that of another enabled **Alias**, Message Analyzer displays a **Duplicated value** message. When this occurs, you will be unable to apply the duplicate **Alias**.

Using the Built-In Alias Collection Items

Message Analyzer provides the following **Aliases** by default in the **My Items** category of your **Aliases** asset collection, which are accessible from the **Aliases** drop-down list in the previously specified locations. You can use these **Aliases** as is, for the convenience of friendly and quick identification of loopback traffic:

- **IPv4 Loopback** — provides a friendly name for the IPv4 address `127.0.0.1`, which is a special IP number that is designated for the software loopback interface on your machine.
- **IPv6 Loopback** — provides a friendly name for the IPv6 address `::1`, which is a special IP number that is designated for the software loopback interface on your machine.

> **NOTE**
>
> A loopback interface is not associated with any hardware, nor is it physically connected to a network. Rather, the interface is a conduit for local application traffic, for example, between a local web server and a SQL server.

You can also modify and save any of the default **Alias** collection items as necessary, as described in Modifying Message Analyzer Aliases. You can also share **Aliases** with others, including the default **Aliases**, by using the

Manage Aliases dialog as described in Managing Message Analyzer Aliases as Shared Items.

Specifying a Default Alias List

If you want to specify a default set of **Alias** names to apply to all message data that you either capture in Live Trace Sessions or load into Message Analyzer through Data Retrieval Sessions, you can simply enable those **Alias** names that you want to automatically apply. You might consider creating a custom **Category** of **Aliases** that will serve as the default set, while manually enabling and applying other **Aliases** as needed. The default set of enabled **Aliases** will persist across Message Analyzer sessions, viewers, and restarts.

> **NOTE**
>
> For a given data field type that is associated with a particular **Alias**, for example, a port or address type, you cannot apply another **Alias** asset with the same **Value** to a message set, although multiple **Alias** assets with the same **Value** can exist in your **Aliases** collection. However, you can only enable and apply **Aliases** with duplicate **Values** one at a time.

Refreshing Views

2 minutes to read

There are several **Alias** operations that you can perform that will trigger a refresh of message data that is displayed in Message Analyzer. In at least one case, the trigger is automatic and in others it is in response to a particular action. These triggers are described in the list that follows:

- **Modifying an applied Alias** — if you edit the **Value** of an **Alias** or its name while it is in an enabled state and you save it, Message Analyzer automatically performs a refresh of the current message set.
- **Enabling/applying an Alias** — whenever you enable a previously disabled **Alias** by selecting its check box, Message Analyzer displays the **Aliases Changed** information bar above the **Analysis Grid** viewer to prompt you to refresh your data views. The information bar also displays the advisory message that is described in earlier topics of this section.
- **Disabling an Alias** — whenever you disable a previously enabled **Alias** by unselecting its check box, Message Analyzer displays the **Aliases Changed** information bar above the **Analysis Grid** viewer to prompt you to refresh your data views.
- **Deleting an Alias** — whenever you delete an **Alias**, Message Analyzer displays the **Aliases Changed** information bar above the **Analysis Grid** viewer to prompt you to refresh your data views.

When any of the previously specified actions occur, Message Analyzer performs a global refresh of the current message set, which includes every session and viewer in the current Message Analyzer instance, including **Layouts** for the **Chart** viewer. The global refresh will also pertain to various Message Analyzer operations that are impacted by application of the **Alias**, such as any applied grouping, sorting, filtering, color rule, or find operation in a viewer where the **Alias** is in use.

Removing the Aliases Changed Information Bar

After the information bar displays, you can remove it in the following ways:

- **Perform a view refresh** — click the **Refresh Views** button to update the current message set with changes that you specified.
- **Remove the changes that you specified** — this applies to modifying, enabling, or disabling an **Alias** only. For example, if you edited an **Alias**, removing the changes would mean to undo the edits. In the case of enabling or disabling an **Alias**, it would mean to undo those actions.
- **Remove the information bar** — click the **X** button on the right side of the information bar to remove it. Note that if you click this button to remove the information bar, it may display again the next time Message Analyzer checks for changes to an **Alias**, at which point the previous update will be detected.

> **NOTE**
>
> The **Aliases Changed** information bar also contains a **Show History/Output Window** icon that enables you to display the **Output Tool Window**. The **Output** window keeps track of various Message Analyzer operations, including installation events, for informational purposes.

Performing Message Analyzer Operations with Aliases

3 minutes to read

Message Analyzer enables you to use **Aliases** in the following common operations that you can employ when analyzing session results:

Using an Alias in a Filter Expression
Sorting Operations with Aliases
Grouping Operations with Aliases

This section briefly describes these capabilities, in addition to the dynamic application of **Aliases** as you are acquiring input message data through a Live Trace Session or a Data Retrieval Session.

Using an Alias in a Filter Expression

Message Analyzer enables you to use an **Alias** in Filter Expressions. Any applied filter that uses an **Alias** should behave in a manner that is similar to any other functionally equivalent filter. However, when creating filters that use **Aliases**, you must consider that the **Alias** name is always a string. For example, if you want to create a view **Filter** that isolates source traffic based on a hypothetical **Alias** named "MyComputer", you might have a Filter Expression that looks similar to the following examples:

```
*Source=="MyComputer"
<moduleName>.Source == “MyComputer”
```

Also note that you can still use the original data type that underlies the **Alias** configuration, for example, an IP address. In this case, the equivalent Filter Expression might look similar to the following:

```
*Source==192.168.1.1
```

You should be aware that because an IP address is a different data type, no quotes are used in the latter Filter Expression.

After you create and save a Filter Expression with the **Edit Filter** dialog — accessible by clicking the **New Filter** item in the **Library** drop-down list on the default Filter panel of the Filtering Toolbar — and that expression uses an **Alias**, as expected the Filter will appear in the centralized Filter Expression **Library** — accessible from any Filter panel on the Filtering Toolbar and on the toolbar above the **Session Filter** text box of the **New Session** dialog. Because any Filter Expression that contains an **Alias** is saved in the centralized Filter **Library**, you can use such an expression as a **Session Filter** or view **Filter**. However, for an **Alias** to function properly in either of these Filter Expressions, the **Alias** must be enabled in the **Aliases** drop-down list, which you can access on the global Message Analyzer toolbar. Also, if you change a Filter Expression that contains an **Alias** in mid-stream during a capture, the changes will immediately take effect.

> **NOTE**
> If you attempt to use the Filter IntelliSense Service to compose a Filter Expression with an **Alias**, you will be unable to find any **Aliases** in the IntelliSense tree because this capability is not yet supported.

Sorting Operations with Aliases

After you create or apply an **Alias**, the data fields in an **Analysis Grid** viewer column for a supported data type

may contain varying types such as integers, strings, or others such as addresses. If you sort such a column by clicking the column header, or by clicking the ascending or descending column arrow, it will be sorted in lexicographical order. For example, sorting results might look like the following:

23.34.51.101
34.56.11.23
ClientComputer
FF02:0:0:0:0:0:1:3

Grouping Operations with Aliases

If you group an **Analysis Grid** viewer column that supports aliasing and an **Alias** is enabled that corresponds with data field value instances that display in that column, Message Analyzer will create a separate group among other group nodes to contain the messages that correspond with the **Alias** name. For example, if you have a **Source** IP address for which you created an **Alias** named "MyComputer" and you perform a grouping operation, by right-clicking the **Source** column in the **Analysis Grid** viewer and selecting the **Group** command, the results will show expandable **Source** nodes (groups) that display the following information:

- The number of messages in parentheses for each node or group.
- The data field value, in this case an IP address or the **Alias** name.

For the hypothetical **Alias** named "MyComputer", the group identifier might look similar to the following:

> Source (110): MyComputer

If you expand this group by clicking the arrow, you will see only the messages that have a "MyComputer" **Alias** in the **Source** column of the **Analysis Grid** viewer.

Dynamically Applying Aliases

If you have a Live Trace Session in progress or if you are in the process of loading static data into Message Analyzer through a Data Retrieval Session, you can create a new **Alias** or enable/disable an existing **Alias** and the results will immediately take effect.

More Information
To learn more about view **Filters**, see Applying and Managing Filters.
To learn more about grouping, see Using the Analysis Grid Group Feature.

Managing Message Analyzer Aliases as Shared Items

3 minutes to read

The **Aliases** drop-down list in the global Message Analyzer **Tools** menu or on the global toolbar contains items that you can share with others. Message Analyzer provides a simple way to expose **Alias** items to others on your team for sharing purposes, or to retrieve **Alias** items that others have shared. You can manage your **Alias** assets in the **Manage Alias** dialog, from where you can **Import** or **Export** an **Alias** asset collection Library containing one or more **Aliases** that you select. Any **Alias** Library that you export is saved as a *.asset file with a name that you specify and any **Alias** Library that you import can only be of the same type. In addition, you can store an **Alias** asset collection Library on a user-configured file share or a feed that you create in the Message Analyzer sharing infrastructure.

Exporting and Importing Alias Libraries

To create an **Alias** asset collection Library for export, you simply select the specific **Aliases** or categories containing one or more **Aliases** in the **Manage Alias** dialog that you want to include in the Library. After you select **Aliases** and click the **Export** button, the **Save Library** dialog displays, in which you can specify **Title**, **Description**, **Author**, and **Organization** information before you **Save** the Library in a chosen location.

To import an **Alias** Library, click the **Import** button in the **Manage Alias** dialog and navigate to the directory location where the **Alias** assets are stored. When you open the Library, you are prompted by the **Select Items to Import** dialog to choose the **Alias** items you want to import. You also have the option to specify the **Category** in which to import the Library items.

Sharing Aliases on a File Share

You can share **Alias** items directly with others by using the **Export** feature in the **Manage Alias** dialog to save one or more **Alias** items as a collection to a designated file share. You can also use the **Import** feature in the same dialog to access **Alias** items that have been shared by others in a similar manner.

Sharing Aliases Through a User Feed

You can also share your **Alias** items through a user feed that you configure in the Message Analyzer Sharing Infrastructure from the **Settings** tab in the **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu. Thereafter, you can use the **Export** feature of the **Manage Alias** dialog to post your **Alias** collection items to the feed so that others can access them on the **Downloads** tab of the **Asset Manager** dialog. As the asset publisher, if you update your existing **Alias** items or add others to your collection, you can make them available to team members or other consumers through the configured feed, where they can view, synchronize with, and download your collection items. However, to enable others to synchronize with item collection *updates*, you will need to perform some manual configuration, as described in Manual Item Update Synchronization.

Downloading Alias Asset Collections from Microsoft

Microsoft also provides a default **Message Analyzer** feed on the **Downloads** tab of the **Asset Manager** dialog that will enable you to download **Alias** asset collections from a Microsoft web service in a future Message Analyzer release. When available, you will be able to synchronize with asset collection updates that are periodically pushed out by the service.

More Information

To learn more about using the common **Manage <*AssetType*>** dialog to share and manage your **Aliases**, see

the Managing User Libraries topic.

To learn more about the Sharing Infrastructure, downloading asset collections, and auto-syncing asset collection updates, see the Sharing Infrastructure and Managing Asset Collection Downloads and Updates topics.

Configuring and Managing Message Analyzer Unions

5 minutes to read

The Message Analyzer **Unions** feature provides a service that can help you overcome the difficulties of analyzing data that derives from different sources with varying naming conventions, in a common environment. If you have an interlaced session containing trace data from multiple sources — for example, .cap, .matp, and text log files — you might also have fields or properties within the session data that are essentially identical, but Message Analyzer does not recognize them as such because they have different names. However, for ease of analysis, you can manually configure Message Analyzer to correlate such fields with a single new entity that you specify as a **Union**. A common example of how Message Analyzer *automatically* performs such a correlation is the manner in which date-time stamps of messages from different supported data sources are treated. Message Analyzer interprets the date-time stamps of such data sources and converts them to a predefined format that displays in a single **TimeStamp** column in the **Analysis Grid** viewer. This enables you to perform operations such as sorting and grouping by time, which are powerful data analysis techniques. However, there can also be many other fields, properties, and annotations from disparate data sources that you can configure Message Analyzer to recognize as identical.

Understanding Data Field Correlation

To facilitate this field correlation capability, Message Analyzer provides the **Unions** feature that enables you to combine one or more message fields, properties, or annotations from different data sources into a single entity that is defined by a **Union** name that you specify. You can then display the named **Union**, which correlates and represents its constituent message field components, as a new data column in the **Analysis Grid** viewer. By configuring a **Union** that maps similar data fields with different names to a single user-defined **Union** name, you can easily correlate equivalent trace data in a message set that combines multiple traces.

> **NOTE**
>
> Message Analyzer also enables **Unions** to support multiple values as "sets" of values for a particular field in the correlation, as described in Creating Unions. In addition, you can use a **Union** that configures a set of values in a **Chart** viewer **Layout**, for example, a set of IP addresses, as in the default **AddressPair Union** that you will find in the **Field Chooser**.

Simplifying Data Analysis with Unions

In the figure that follows, you will see messages that were captured at two different interfaces by two different data sources, a SambaSysLog and a live Wireshark trace. The scenario represented by this illustration involves troubleshooting SMB file access issues. To facilitate analysis, data from both sources is loaded together into Message Analyzer through a Data Retrieval Session and interlaced as if they were one trace, so that the log messages and associated SMB live capture messages can be assessed together.

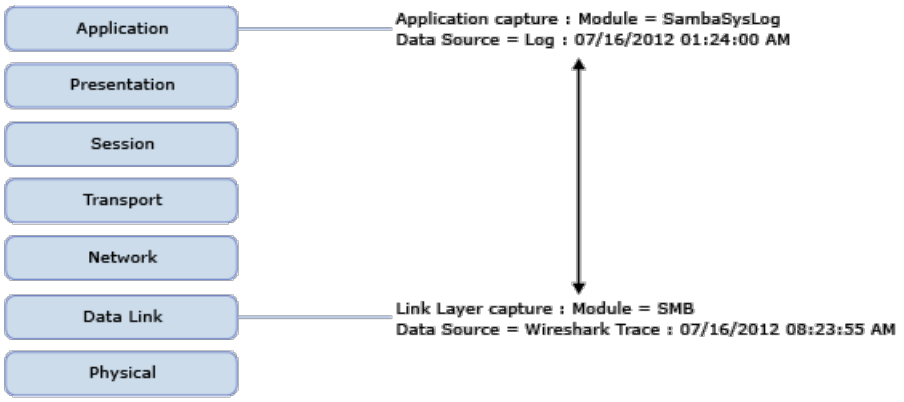

Module	Data Source	Command.smb_cmd	Command	SMBCommand2
SambaSysLog	log - 07/16/2012 01:24:00 AM	117		117
SambaSysLog	log - 07/16/2012 01:24:00 AM	50		50
SambaSysLog	log - 07/16/2012 01:24:00 AM	50		50
SambaSysLog	log - 07/16/2012 01:24:00 AM	114		114
SMB	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_NEGOTIATE(114)	114
SMB	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_SESSION_SETUP_ANDX(115)	115
SMB	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_SESSION_SETUP_ANDX(115)	115
SMB	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_TREE_CONNECT_ANDX(117)	117
SambaSysLog	log - 07/16/2012 01:24:00 AM	117		117
DFSC	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_TRANSACTION2(50)	50
SMB	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_TRANSACTION2(50)	50
SMB	Wireshartktrace -0 7/16/2012 08:23:55 AM		SMB_COM_TREE_CONNECT_ANDX(117)	117
SMB	Wireshartktrace - 07/16/2012 08:23:55 AM		SMB_COM_TREE_DISCONNECT(113)	113
SambaSysLog	log - 07/16/2012 01:24:00 AM	50		50
SambaSysLog	log - 07/16/2012 01:24:00 AM	50		50
SambaSysLog	log - 07/16/2012 01:24:00 AM	50		50
SambaSysLog	log - 07/16/2012 01:24:00 AM	50		50
SambaSysLog	log - 07/16/2012 01:24:00 AM	50		50
SMB	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_TRANSACTION2(50)	50
SMB	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_TRANSACTION2(50)	50
SMB	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_TRANSACTION2(50)	50
SMB	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_TRANSACTION2(50)	50
SMB	Wiresharktrace - 07/16/2012 08:23:55 AM		SMB_COM_TRANSACTION2(50)	50

Figure 63: SMB fields Union example

In the figure, note that the data from the indicated sources contains command values that are virtually equivalent; however, those values are contained under a different column name for each data source (`Command.smb_cmd` and `Command`). Message Analyzer will not recognize such equivalence until you configure a **Union** that correlates the disparate field names that have the same meaning. As illustrated in the figure, this would mean taking the **Union** of the `Command.smb_cmd` and `Command` fields and creating a new entity named `SMBCommand2` . Thereafter, the unified field correlation displays as a single set of values in the `SMBCommand2` column, which you must add to the **Analysis Grid** viewer by using the **Field Chooser Tool Window**. You can even remove the existing data columns for the `Command.smb_cmd` and `Command` fields and the **Union** will continue to display the indicated value set in the `SMBCommand2` column.

Using the Built-In Union Assets

Message Analyzer has several built-in **Unions** of which you can take advantage. You can find these built-in **Unions** in the **Field Chooser** window under the **Unions** node. You can add one or more of these **Unions** as new columns in the **Analysis Grid**; however, this will be useful only if it is appropriate for the data you are analyzing. For example, if you are working with an SMB trace and a SambaSysLog that intrinsically have different field naming conventions, you might add the **SMBCommand** union as a new column to merge the **SMB.Command** field of the trace and the **SambaSysLog.smb_command.command.smb_com** field of the log. You can add this new column to the **Analysis Grid** viewer by right-clicking the **Union** in **Field Chooser** and selecting the **Add as Column** command in the context menu that appears.

You can also review the configuration of the built-in **Unions** in the **Manage Unions** dialog, by right-clicking a

Union and selecting the **Create a Copy** command in the context menu that appears. This action displays the **Edit Union** dialog in which you can view the built-in **Union** field configurations and other settings. The **Manage Unions** dialog is accessible by clicking **Unions** on the global Message Analyzer **Tools** menu and then selecting the **Manage Unions** item in the drop-down list.

The built-in **Unions** are contained in the categories specified below and are described as follows:

- **File Sharing** category
  - **SMBTID** — correlates the following field names for SMB tree IDs from an SMB trace file and a SambaSys log file into the specified **Union** name:
    - **SMB.SmbHeader.Tid**
    - **SMB2.SMB2Request.Header.TreeId**
    - **SambaSysLog.Smb_command.command.smb_tid**
  - **SMBMID**— correlates the following field names for SMB message IDs from an SMB trace file and a SambaSys log file into the specified **Union** name:
    - **SambaSysLog.Smb_command.command.smb_mid**
    - **SMB2.ReadRequest.Header.MessageId**
    - **SMB.SmbHeader.Mid**
  - **SMBCommand** — correlates the following field names from an SMB trace file and a SambaSys log file into the **Union** named **SMBCommand**:
    - **SMB.Command**
    - **SambaSysLog.Smb_command.command.smb_com**
- **SharePoint** category
  - **SharePointCorrelation** — correlates the following field names in a SharePoint Unified Logging Service (ULS) text log file (*.log) and a trace file that captured SharePoint network traffic, respectively, into the specified **Union** name:
    - **ULS.EventHeader.Correlation**
    - **HTTP.Request.Headers.SPRequestGuid**
- **Common** category
  - **PID** — correlates the following field names in an event log (*.etl) and a set of trace results, for example, from a *.matp file into the specified **Union** name:
    - **ProcessId**
    - **Etw.EtwProviderMsg.EventRecord.Header.ProcessId**
  - **AddressPair** — correlates field names into a set of address pairs that display in a single data column in the **Analysis Grid** viewer, which includes source and destination address pairs for each the following:

    Note This **Union** enables sets by specifying the **Set of multiple values** option.
    - **Ethernet.Frame.Source**, **Ethernet.Frame.Destination**
    - **IPv4.Datagram.Source**, **IPv4.Datagram.Destination**

- **IPv6.Datagram.Source**, **IPv6.Datagram.Destination**

Using Unions in Message Analyzer Operations

Message Analyzer also enables the ubiquitous use of **Unions** with various operations, viewers, and tools that you might typically employ during data analysis processes, as follows:

- Grouping
- Sorting
- Filtering
- Pattern matching
- Analyzing data in different viewers, including **Charts**
- Using **Tool Windows**, such as the **Field Chooser**

For further details about using **Unions** in the manner described in the preceding list, see Performing Message Analyzer Operations with Unions.

What You Will Learn
In the following topics of this section, you will learn how to create, modify, manage, and share **Unions**, in addition to performing Message Analyzer operations with them:
Creating Unions
Modifying Unions
Refreshing Data Views Containing Unions
Performing Message Analyzer Operations with Unions
Managing Unions as Shared Items

Creating Unions

8 minutes to read

Message Analyzer enables you to create a **Union** that correlates two or more fields, properties, or annotations, as described in Configuring and Managing Message Analyzer Unions. Creating a **Union** can help streamline data analysis and can also simplify what might otherwise be cryptic field names that are difficult to work with. You can create a new **Union** by selecting the **New Union** item from the **Unions** drop-down list in the global Message Analyzer **Tools** menu or by clicking the **Unions** button on the global Message Analyzer toolbar. Making such a selection causes the **Edit Union** dialog to display, from where you can provide the input configuration data that is required to create a new **Union**.

Go To Procedure

To go directly to a procedure that creates a **Union**, see Create a Union of Two Data Fields. However, you are advised to review the information contained in this section before doing so.

Configuring a Union

To configure a **Union** of two or more data fields, properties, or annotations, use the following controls in the **Edit Union** dialog:

- **Name** — in this text box, specify a **Union** name that is appropriate for your environment. However, note that you cannot prefix a **Union** name with a number; otherwise, you can follow your own naming conventions. Note that Message Analyzer permits you to specify a **Union** name that is similar to one of its child (constituent) field names. For example, a **Union** could be named "command", even if one of its child field names is "SMB.command". This is because the result of such a naming, <*UnionName*>.command, is unique and would not collide with the child field name.

> **NOTE**
>
> If the **Union** name that you specify already exists, Message Analyzer displays a message box with the following prompt: "**The name has been used, please use a different name**"

- **Type** — consists of a label that indicates the **Union** data type. It is likely that the fields, properties, or annotations that you combine into a **Union** will be of different data types. When this is the case, Message Analyzer performs a fundamental conversion to an appropriate type, as described in Supporting Type Conversions for Union Fields, to ensure that an appropriate data type is generated for your **Union**.

> **NOTE**
>
> If you add or remove any field, property, or annotation in the **Edit Union** dialog, the **Type** label may update to reflect a new **Union** data type, depending on the field types you are adding to or removing from the **Union** configuration.

- **Category** — in this combo box, specify the name of the **Category** in which to place your new **Union**. You have the option to either create a new **Category**, which is retained as an editable drop-down list item for future selection, or you can choose an existing **Category** from the drop-down. Be sure to create or choose a **Category** that is meaningful for your environment. All new categories that you specify become subcategories that appear under the default **My Items** top-level category in the **Unions** drop-down list. However, if you leave the **Category** combo box blank, the **Union** will be stored directly under the **My**

Items category.

- **Add** — click this button to display the **Field Chooser Tool Window**, from where you can select the fields, properties, or annotations that you want to combine as a **Union** and create a correlation of fields with equivalent meaning. After you locate each field in the **Field Chooser** window that you want to add to the **Union**, either double-click the name of the field, property, or annotation; or highlight it and click the **Select** button in the dialog to add it to the **Union** configuration. Each field entity that you specify in this manner is added to the **Select fields to include** list box. To create a **Union**, you must have at least two fields, properties, or annotations.

  If you want to get the most out of this feature when you are combining fields, properties, or annotations into a **Union**, you should carefully consider the following suggestions:

  - The added entities essentially have the same functional meaning, otherwise the resulting **Union** is likely to be irrelevant.
  - The added entities are optionally of the same data type, to minimize memory consumption (see Supporting Type Conversions for Union Fields).
  - The added entities come from *different* data sources rather than the *same* data source or module, as this is the intended use of a **Union**.

    > **IMPORTANT**
    >
    > If you want to create a **Union** of two fields, one from a trace file and one from a log file, they will need to be file types that Message Analyzer supports and the log file will have to be parsed so that a module node for the log will appear in the **Field Chooser** window, from where you can select the fields of the log. To parse text logs, Message Analyzer requires an OPN configuration file. Several configuration files are provided with Message Analyzer by default to parse various types of text logs, such as SambaSysLogs, Netlogon logs, IIS logs, Cluster logs, and others.
    >
    > To initiate the parsing process, you will need to specify a configuration file in the **Text Log Configuration** drop-down list in the **New Session** dialog for a Data Retrieval Session (**Files** tab) and then click the **Start** button in the dialog. Thereafter, the name of the log will appear as a node in **Field Chooser**, from where you can select any of the fields that Message Analyzer parsed for a particular log type whenever you are creating a **Union**.

Without considering data sources, you might specify arbitrary fields that hold unrelated data, in which case Message Analyzer will only use the first field that you configured in the **Union** — note that this can have unexpected results.

- **Return value as** options — consist of the following:
  - **Single value** — this option causes the **Union** to display a single value, even if the correlation results in multiple values for a particular field. When multiple values exist, Message Analyzer picks a value for the **Union** to display, based on the first field that you configured in the correlation.
  - **Set of multiple values** — this option causes the **Union** to display a set of values when the correlation results in more than one value for a field. For example, Message Analyzer might display a set of IP Addresses as "set (192.168.1.1, 192.168.1.2)" in the **Analysis Grid** column that displays the correlated **Union** values, assuming you add the named **Union** as a new column in the **Analysis Grid** viewer with the **Field Chooser** window.
- **Remove** — click this button to remove an existing field that you highlight in the **Select fields to include** list box. Note that you can remove fields only one at a time.
- **Save** — click this button to save the **Union** configuration. At this time, a new **Union** is created and added to the **Unions** drop-down list in the **Category** that you specified.

Note Your new **Union** will also be added to the root **Unions** node in the **Field Chooser** window after a Message Analyzer restart. Note that you will use the **Field Chooser** window to add your new **Union** as a column in the **Analysis Grid** viewer. In addition, the new **Union** becomes an asset that you can share with others through the Message Analyzer sharing infrastructure, as described in Managing Unions as Shared Items.

Supporting Type Conversions for Union Fields

When you combine two or more fields, properties, or annotations into a **Union**, it is often the case that such entities are of different data types. For instance, you might be combining `byte` and `int` data types in a **Union**, which have incompatible range values and require a type conversion for use in the **Union**. The following describes how Message Analyzer treats compatible and incompatible field types, with respect to type conversions:

- **No type conversion** — compatible field types do not require any type conversion, such as those fields that are of the *same* data type. This results in the **Union** taking on the same data type as the correlated fields. For example, a **Union** with two correlated fields of type `ushort`, will display **ushort** in the **Type** label of the **Edit Union** dialog.

- **Implicit type conversion** — Message Analyzer can accommodate for incompatible field types through the use of a transitive implicit type conversion, if an appropriate implicit conversion path exists for one or more of the correlated fields. Basically, a transitive conversion says that if type A can be converted to type B and type B can be converted to type C, then type A can be converted to type C as well. Message Analyzer performs such a transitive type conversion in order to detect the appropriate data type of minimum sufficient value range that can represent all the types in the **Union**. This minimalistic approach ensures that the **Union** consumes the least amount of memory possible. In the case of a **Union** with fields of type `byte` and `int`, the `int` is the minimum type that has a suitable value range to handle the range of values that can occur in fields of these types in the **Union**. Therefore, the **Union** takes `int` as its type, which then displays in the **Type** label in the **Edit Union** dialog.

More Information
The [OPN Programming Guide] (https://download.microsoft.com/download/3/E/8/3E845130-349C-4EFC-B634-C7DBD4614 0B7/OPN%20Programming%20Guide%20v4.4.docx) contains a type conversion table in section 2.1.7 that enables you to map the conversion path taken by Message Analyzer in the previously mentioned type conversion. For example, the type `byte` can implicitly convert to type `short`, and a `short` can implicitly convert to an `int`, therefore, type `byte` can convert directly to type `int`.

- **Base type conversion** — incompatible field types would remain incompatible if there were no implicit conversions to a common type available. Therefore, when this is the case, Message Analyzer converts all field types to a common base type called `any`. The `any` type is sufficient to handle the value range of any other data type. The **Union** would then display the `any` type in the **Type** label of the **Edit Union** dialog.

Performing Other Operations on Unions

After you successfully create one or more **Unions**, you can perform the following operation on any **Union** that exists in the **Manage Unions** dialog, as described in Modifying Unions:

- **Create a Copy** — this is the only command available for built-in **Unions**, given that built-in **Union** assets cannot be edited.

You can perform the above operation in addition to the following on any **Union** that exists in the **My Items** category of the **Manage Unions** dialog. Note that these commands are not available for the built-in **Unions**:

- **Edit**

- **Delete**

Modifying Unions

5 minutes to read

After you configure one or more **Unions** according to the instructions outlined in Creating Unions, you have the option to modify, delete, or create a copy of any **Union** you have created. To do this, you can use the Message Analyzer commands that appear in the **My Items** category of the **Manage Unions** dialog. These commands are available from the context menu that displays in the **Manage Unions** dialog whenever you right-click a **Union** in the **My Items** category of the dialog. You can access the **Manage Unions** dialog by clicking the **Unions** item on the global Message Analyzer **Tools** menu and selecting the **Manage Unions** item. To edit a **Union** in the **My Items** category, you must right-click the **Union** you want to edit and then select the **Edit** command from the context menu that appears. This action displays the **Edit Union** dialog from where you can modify the **Union**. The other commands in this context menu are **Create a Copy** and **Delete**, as described in the next section.

Commands for Modifying Unions

After you right-click a **Union** in the **My Items** category, you can make use of the following commands to modify it.

- **Edit** — click this command to open the **Edit Union** dialog, in which you can use the same controls to modify a **Union** that you use when creating a new **Union**, as described in Creating Unions. The modifications that you can perform or effect from the **Edit Union** dialog are listed below.
  - **Edit a Union name** — edit a **Union** name by entering your changes in the **Name** text box.
  - **Change a Union category** — select a different category for the **Union** from the **Category** drop-down list, or enter a new category name in the **Category** combo box.
  - **Add fields to a Union** — click the **Add** button to open the **Field Chooser Tool Window**, from where you can add one or more fields, properties, or annotations to the **Union**, one at a time. Note that this action can result in a **Type** change.
  - **Remove fields from a Union** — highlight a field that you want to remove and click the **Remove** button to exclude it from the **Union**. This too can result in a **Type** change.
  - **Change message data** — if you add or remove fields, properties, or annotations in a **Union** that does not cause a corresponding **Type** change, and the **Union Name** also remains unchanged, this can result in an immediate update to the underlying message data in any active data viewer that uses the **Union**. This might include currently active viewers such as the **Analysis Grid** or **Chart** viewer. In addition, if a sort, group, or other operation is in progress when such a data change occurs, all active viewers will restart those operations. Lastly, you might observe that the **Change Union** dialog does not display in this case because a message data change is immediate and does not require a Restarting Message Analyzer.

> **NOTE**
>
> If there is a change to a **Union** field, property, or annotation and there is also a **Type** or **Name** change involved, then the **Change Union** dialog displays to indicate the type of change that will be applied after Message Analyzer is restarted. The dialog also indicates that you should remove the **Union** from any Message Analyzer features that are configured to use it, for example, a Filter Expression or a Chart **Layout**. In addition, the dialog prompts you to continue or not, which you decide by clicking either the **Yes** or **No** button, respectively. Note that the previously existing **Union** name remains active and that the **Union** still functions normally until Message Analyzer is restarted.

- **Create a Copy** — click this command to open the **Edit Union** dialog that is automatically populated with a replica of the existing **Union** configuration. You can then modify the copy and **Save** it under a new **Union** name. By providing a **Union** template, this command enables you to quickly create a new **Union** based on similar data from an existing one. When working with a copy of an existing **Union**, the following guidelines apply:
  - **Saving a Union copy that matches an existing name** — if you try to **Save** a copy of an existing **Union** without changing its **Name** value, Message Analyzer displays a message to indicate that the **Union** name is already in use. Message Analyzer does not allow you to save such a **Union** because all **Union** names must be unique.
  - **Saving an edited Union copy** — if you modify a copied **Union** and **Save** it with a **Union** name that is different than any other **Union** name, Message Analyzer adds it to the **Unions** list in the **Manage Unions** dialog, whether or not a **Type** change has also occurred. The **Union** is also added to the **Field Chooser** window under the **Unions** node. Note that changes you make to a copied **Union** can include a **Type** change, for example, by adding or removing a field that causes a **Type** change.
- **Delete** — click this command to delete a **Union**.

  The deleted **Union** is then removed from the **Unions** list in the **Manage Union** dialog, as it is no longer an editable or shareable asset. Note that any **Analysis Grid** viewer column that contains the **Union** name continues to persist. However, after the **Union** is removed at the next Message Analyzer restart, it will no longer be available from **Field Chooser** and therefore you will be unable to add it as a new data column in the **Analysis Grid** viewer. Until such time, you can continue to display and use the **Union** as needed.

Restarting Message Analyzer

When you restart Message Analyzer, any changes that are queued up for one or more **Unions** will be applied. Note that any **Union** you delete in the **My Items** category of the **Manage Unions** dialog will be removed from the root **Unions** node in the **Field Chooser** window following a Message Analyzer restart.

Consumers of Changed Unions

If a changed **Union** is called by a Message Analyzer feature that uses the **Union** (for example, a Chart **Layout** formula or Filter Expression), Message Analyzer handles these situations within the context of existing functionality. For instance, if you attempt to create a Filter Expression that utilizes a **Union** that has had a **Name** change, or that you removed with the **Delete** command, an automatic compilation check that occurs before the filter is applied to a message set might display the **Compile Query Error** dialog. This dialog indicates that the Filter Expression failed to compile because the **Union** entity could not be found. Note that a similar process takes place for **Pattern Expressions**, which are also subject to a compilation check.

> **IMPORTANT**
> If you have any Message Analyzer feature that uses a **Union** — for example, any type of **Layout** or Filter Expression — that has undergone a change in **Name**, **Type**, or is pending deletion, you are advised to manually remove such a **Union** from those features, as described in Pending Changes.

See Also

Creating Unions
Refreshing Data Views Containing Unions

Refreshing Data Views Containing Unions

2 minutes to read

This section briefly describes how Message Analyzer data and operations are refreshed in response to **Union** modifications.

Refreshing Message Data

If you change a **Union** field, property, or annotation but the **Type** does not change, and you do not modify the **Name** value of the **Union**, then an automatic refresh of message data is performed for all open data viewers in which the **Union** is active. For example, you might have a **Union** displaying in the **Analysis Grid** viewer as a result of loading a view **Layout** that contains a column that is named by the **Union**. Moreover, you might have a **Layout** for the **Chart** viewer that uses the **Union** in a visualizer component where you have configured a formula to perform a particular calculation for the data display. In these cases, a refresh of message data could involve an update to the displayed set of **Union** values. The previously indicated type of change to a **Union** is the only one that can result in an automatic refresh of message data in all data viewers and visualizers that are currently open and using the **Union**. In all other cases, there is no automatic refresh of data, as described in Pending Changes.

Refreshing Operations in Progress

If Message Analyzer is sorting, grouping, filtering, or performing some other operation on a message set when a **Union** data change occurs, all applicable data viewers that use the **Union** will restart those operations.

Pending Changes

If you modify a **Union** by changing the **Name**, or if the **Type** changes because of adding or removing a **Union** field, property, or annotation, or if you delete the **Union**, then message data in all open data viewers that use the **Union** is *not* refreshed. Instead, the **Union** is internally marked for either the edited state or deleted state and awaits a Message Analyzer restart before the pending changes are applied.

At this time, you are advised to manually remove such a **Union** from all Message Analyzer features in which it is currently used. For example, if there was a **Name** change for the **Union**, after a restart it will no longer exist under its previous **Name** in any feature that formerly used it. Therefore any Filter Expressions, view **Layouts**, Chart **Layout** formulas, **Color Rules**, and so on, that use such a **Union** will be impacted. Note that even the display of **Aliases** can be affected. For example, a **Union** *value* that displays in an **Analysis Grid** column that is named by the **Union**, could be an **Alias** due to the underlying alias **Value** (as in the **Alias Editor** dialog), and **Union** modifications can change the way an **Alias** displays.

Performing Message Analyzer Operations with Unions

4 minutes to read

After you create a **Union**, as described in Creating Unions, it is a new entity that can function normally in various Message Analyzer operations. However, to facilitate the use of a **Union** in most of the operations described in this section, you will first need to add a column to the **Analysis Grid** viewer based on the **Union** name. To do this, open the **Field Chooser Tool Window** by clicking the **Add Columns** button on the **Analysis Grid** viewer toolbar. You can navigate to the **Union** of interest by clicking the root **Unions** node in the **Field Chooser** window to expose the **Unions** that it contains. You can then double-click the **Union** name to add it as an **Analysis Grid** viewer column. Thereafter, you can observe that such a **Union** functions as expected in the following Message Analyzer operations:

- **Sorting** — to sort a **Union** column in the orders indicated below, click the column header as follows:
  - **Ascending** — the first click on the **Union** column header sorts the **Union** values in ascending order, as indicated by an up arrow that appears on the column header, providing that the column header did not already display an up or down arrow.
  - **Descending** — the second click on the **Union** column header sorts the **Union** values in descending order, as indicated by a down arrow that appears on the column header.
  - **Unsorted** — the third click on the **Union** column header returns the messages to the order in which they were initially captured by Message Analyzer.

    NOTE

    If you continue to click the **Union** column header after the initial three clicks indicated in this list, the column will be sorted in round-robin fashion, with ascending, descending, and unsorted results repeating over and over, in that order.

- **Grouping** — after you add the **Union** as a column in the **Analysis Grid** viewer as previously indicated, you can perform a Grouping operation, just as you do with any other **Analysis Grid** viewer column that contains a related set of values. In this case, the Grouping operation results in creating groups of data that are organized according to the different values that exist in the column with the **Union** name. You can continue with additional Grouping operations on other **Analysis Grid** viewer columns that will typically result in nesting additional child data groups within the original parent **Union** group. By pivoting on various combinations of column data, you can create a variety of unique data analysis perspectives.

NOTE

Because **Unions** now support sets, that is, multiple values for a particular field correlation, a **Grouping** operation will reflect the set values.

- **Filtering** — after you create and save a **Union**, it is available for use in a Filter Expression. The simplest way to create a Filter Expression with a **Union** is to right-click any field value in the **Union** column and select the **Add '<*unionName*>' to Filter** item, where '<unionName>' is a placeholder for an actual **Union** name. This might be the quickest way to create a valid Filter Expression with a **Union**, because Message Analyzer automatically passes the associated module into the Filter Expression and provides the

correct syntax for the filter, as expressed in the following format: `module.unionName==someValue`. The Filter Expression then appears in the text box of the default Filter panel on the Filtering toolbar in the indicated format.

After you create a filter in this manner, you can add it to your centralized Filter Expression **Library** by selecting the **New Filter** item in the **Library** drop-down list in the default Filter panel on the Filtering toolbar. The text of the Filter Expression containing a **Union** is automatically passed to the **Edit Filter** dialog that displays. After you save the Filter Expression that contains your **Union**, you can use it in any Message Analyzer feature that has access to the centralized **Library**, which includes a **Session Filter** or view **Filter** that you can specify when you run a new session or when you are analyzing session results, respectively.

- **Pattern matching** — after you create and save a **Union**, it is available for building a Pattern Expression. For example, you can access the **Union** in the **Field Chooser Tool Window**, that you typically use to locate fields when creating a Pattern Expression. You can treat the **Union** the same way you do any other field, property, or annotation when using the **Pattern Editor** to build a Pattern Expression, although **Union** behavior in the Pattern Expression might be different than other fields, depending on how you construct the expression.
- **Color Rules** — after you create and save a **Union**, it is available for use in Filter Expressions that you can add to a **Color Rule** from the centralized Filter Expression **Library**. In the context of **Color Rule** functionality, the Filter Expression that contains a **Union** will behave in a manner that is similar to any Filter Expression that contains a comparable type. This is also true of other features that can use a Filter Expression, for example, a **Chart**.
- **Layouts** — after you create a **Union**, you can use the **Field Chooser** window to add it to the **Analysis Grid** viewer as a new column. Thereafter, you can save the current **Analysis Grid** column configuration as a new **Layout** that contains the **Union** column, which you can apply any time you want to recover that particular column configuration.

See Also

Using the Analysis Grid Group Feature
Filtering Live Trace Session Results
Pattern Match Viewer
Creating and Modifying Color Rules
Applying and Managing Analysis Grid Viewer Layouts

Managing Unions as Shared Items

5 minutes to read

The **Manage Unions** dialog contains **Union** items that you can share with others. Message Analyzer provides a simple way to expose **Union** items to others on your team for sharing purposes, or to retrieve **Union** items that others have shared. You can manage your **Union** assets in the **Manage Union** dialog, which is accessible as an item in the **Unions** drop-down list on the global Message Analyzer **Tools** menu. From this dialog, you can **Import** or **Export** a **Union** asset collection Library containing one or more **Unions** that you select. Any **Union** Library that you export is saved as a *.asset file with a name that you specify, and any **Union** Library that you import can only be of the same type. In addition, you can store a **Union** Library on a user-configured file share or a feed that you create in the Message Analyzer sharing infrastructure.

Working with Union Libraries

To create a **Union** asset collection Library for export, you simply select the specific **Unions** or categories containing one or more **Unions** in the **Manage Union** dialog that you want to include in the Library. After you select **Unions** and click the **Export** button, the **Save Library** dialog displays, in which you can specify **Title**, **Description**, **Author**, and **Organization** information before you **Save** the Library in a chosen location.

To import a **Union** Library, click the **Import** button in the **Manage Union** dialog and navigate to the directory location where the **Union** asset file is stored. When you open the file, you are prompted by the **Select Items to Import** dialog to choose the **Union** items you want to import. You also have the option to specify the **Category** in which to import the Library items.

The **Manage Unions** dialog also enables you to modify any **Union** that displays in the **My Items** category of the dialog list. You can access the modification commands by right-clicking any **Union** in the dialog to display a context menu that contains the **Edit**, **Create a Copy**, and **Delete** commands. If you **Delete** a **Union** in the **My Items** category of the **Manage Unions** dialog, obviously that **Union** will not be available to include in an export Library. If you decide to **Edit** a **Union** in the **My Items** category from the dialog, or **Create a Copy** and edit the **Union** copy, and you save it, then only the edited version of the **Union** is available to include in a Library export. If a user that is running a different Message Analyzer instance imports the Library, the **Union** assets that it contains will be new entities on their computer and will take effect immediately, although it is still possible that there could be a collision with an existing **Union** name. However, any modifications that you made to a **Union** in the current Message Analyzer instance take effect only after a Message Analyzer restart.

Sharing Unions on a File Share

You can share **Union** items directly with others by using the **Export** feature in the **Manage Union** dialog to save one or more **Union** items to a designated file share. You can also use the **Import** feature in the same dialog to access **Union** items that have been shared by others in a similar manner.

Sharing Unions Through a User Feed

You can also share your **Union** items through a user feed that you configure in the Message Analyzer Sharing Infrastructure from the **Settings** tab of the **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu. You can specify a feed name and location in the **Add Feed Location** dialog, which is accessible by clicking the **Add New Feed** button on the **Settings** tab. Thereafter, you can use the **Export** feature of the **Manage Union** dialog to post your Library of **Union** assets to the feed so that others can access them on the Message Analyzer **Downloads** tab of the **Asset Manager** dialog. As the asset publisher, if you update your existing **Union** items or add others to your collection, you can make them available to team members or other

consumers through the configured feed, where they can view, synchronize with, and download your collection items. However, to enable others to synchronize with your item collection updates, you will need to perform some manual configuration, as described in Manual Item Update Synchronization.

Accessing the Union Assets Collection

Unions are contained in an asset package named **Message Analyzer Correlations** on the **Downloads** tab of the **Asset Manager** dialog. You can either auto sync the package for periodic updates that occur automatically, or you can download the package from the **Downloads** tab of the **Asset Manager**, at which time the auto-sync status is canceled. However, after you download, you still have the option to auto-sync again later to synchronize your Message Analyzer installation for future updates to this asset package.

> **IMPORTANT**
>
> The **Message Analyzer Correlations** asset collection is installed by default with Message Analyzer, so it is unnecessary to download the collection whenever you start Message Analyzer. But if it is the first time you have started Message Analyzer, you are presented with a **Welcome** dialog that provides you with the choice to opt in or out of automatic updates. If you choose to opt in to auto-syncing updates, then all Message Analyzer asset collections are automatically set to the auto-sync state, including the **Message Analyzer Correlations** asset collection, and no further action is required.

However, if you opted out, you still have the option to automatically receive periodic collection updates later by setting the **Offline** mode to **Online** on the **Downloads** tab of the **Asset Manager** and clicking the **Sync All Displayed Items** button, which auto-syncs all asset collections, or you can set individual collections to the auto-sync state on the **Downloads** tab as you require them. To do this, click the download icon to the right of the collection on the **Downloads** tab and select the **Automatically sync item collection updates when available** option in the **Item Download Options** dialog.

More Information

To learn more about using the common **Manage <*AssetType*>** dialog to share and manage your **Unions**, see the Managing User Libraries topic.

To learn more about the Sharing Infrastructure, downloading asset collections, and auto-syncing asset collection updates, see the Sharing Infrastructure and Managing Asset Collection Downloads and Updates topics.

Viewing Process Name Data

2 minutes to read

The capability to view process names in message data captured by any ETW trace provider is now native to Message Analyzer, although detection of process names is currently not guaranteed for incoming messages. This means that you can add the **ProcessName** field (from the **Global Properties** node of **Field Chooser**) as a new **Analysis Grid** viewer column and view process name data across a set of trace results. If you also add a **Network** field column from the **IPv4** node in **Field Chooser**, you can correlate the IP conversations with which the process names are associated. The input file types in which you can view process name data include .matp, .etl, .evtx, and .cap files.

Using the ProcessName Property

If you want to isolate the messages that were captured by Message Analyzer for each process, you can execute the **Group** command on the **ProcessName** column of the **Analysis Grid** viewer to separate the trace messages into groups of **ProcessName** nodes, where each node contains all the messages associated with a particular process name. You can create this grouped display configuration by right-clicking the **ProcessName** column header and then selecting the **Group** command. If you also added the **Network** field as a new **Analysis Grid** viewer column, as suggested earlier, you can similarly execute the **Group** command on this column to correlate the associated network conversations with process names.

Alternatively, you could simply display the **Process Name and Conversations** view **Layout** for the **Analysis Grid** from the **Layout** drop-down list on the **Analysis Grid** viewer toolbar to view similar data. However, note that this **Layout** also adds a **Transport** group that exposes the ports that carried the network conversations. In any case, the data can tell you very quickly which processes are consuming the most bandwidth and can also help you isolate any process (and supporting messages) that you may already suspect is causing a problem.

For incoming messages, Message Analyzer does not guarantee the display of process names. In this case, Message Analyzer should display the ETW **ProcessID** value in the **ProcessName** column of the **Analysis Grid** viewer.

Layouts Containing the ProcessName Field
The **ProcessName** property is used in the following data viewer **Layouts**:

- **Grouping** viewer — uses the **ProcessName** and **ProcessId** properties in this **Layout**:
  - **Process Name and Conversations** — this **Layout** (left side of the user interface) simulates the **Network Conversation** tree in Microsoft Network Monitor, as shown in the figure that follows.

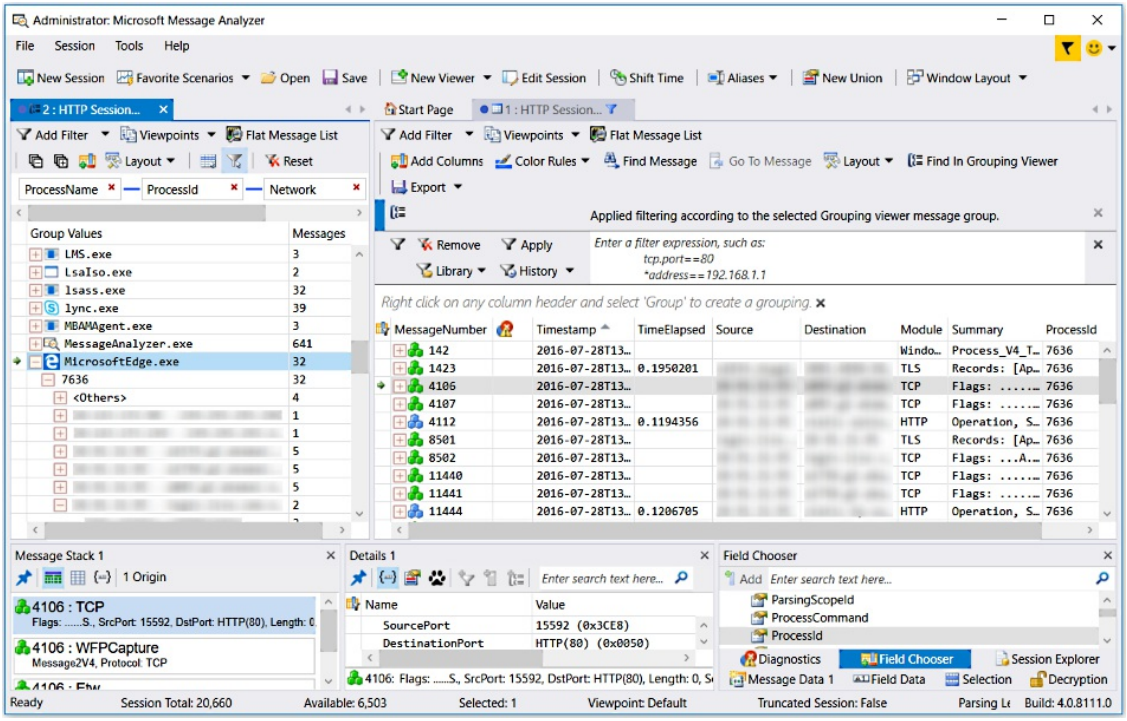

Figure 64: Grouping Viewer ProcessName node selection driving the Analysis Grid viewer

- **Analysis Grid** viewer — uses the **ProcessName** property in these **Layouts**:
  - **Process Name and Conversations**
  - **Network Monitor**
  - **Process View**

See Also

Analysis Grid Viewer
Grouping Viewer
Working With Message Analyzer Profiles

Tool Windows

2 minutes to read

This section describes various interactive **Tool Windows** that are provided with Message Analyzer to display additional message details such as field information, message stack, hexadecimal, diagnosis, and decryption data. Some **Tool Windows** support data analysis capabilities such as the application of **Viewpoints**, view **Filters**, and selection of messages. Others enable you to configure annotations such as comments and bookmarks, or to add new data columns to the **Analysis Grid** viewer based on selected message fields.

Certain **Tool Windows** provide a multi-instance capability that enables you to display up to a maximum of four windows of the same type. This feature enhances your data analysis process because it enables you to compare different message data across multiple instances of the same window type. The **Tool Windows** that support this capability include the **Details**, **Message Data**, and **Message Stack** windows.

The **Analysis Grid** viewer and some common **Tool Windows** are shown in the figure that follows. Note that a number of **Tool Windows** are grouped in the lower right sector of the Message Analyzer user interface, where you can display any chosen **Tool Window** by selecting its tab.

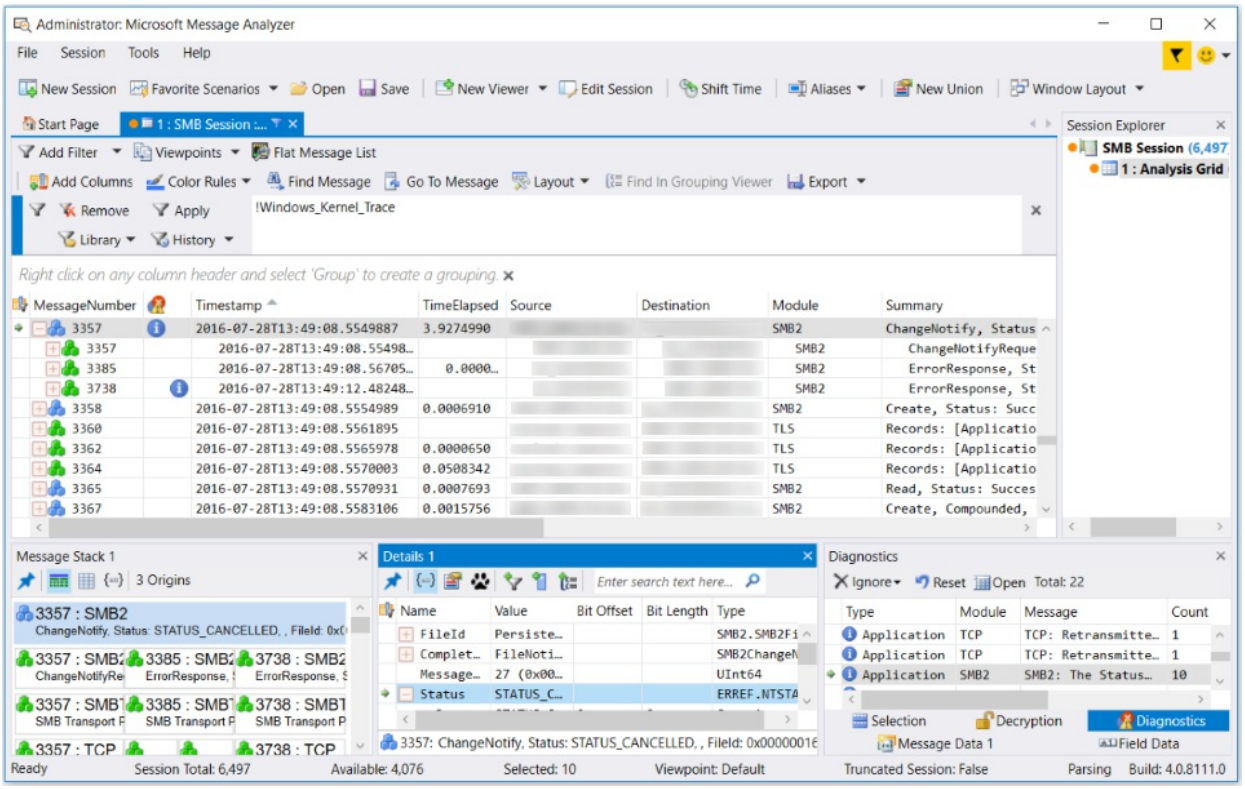

Figure 65: Analysis Grid viewer and Common Tool Windows

The ability to maintain the context of displayed data through multiple message selection is described in further detail in the associated **Tool Window** topics of this section.

Some **Tool Windows** are *message-specific* while others are *session-specific*, meaning that they respond to message or session selection, respectively, and display data that is associated with the in-focus message or session. An example of a message-specific **Tool Window** is the **Message Stack** window and an example of a session-specific **Tool Window** is the **Diagnostics** window, which are both described within the topics of this section along with other **Tool Window** types:

Message-Specific Windows

Session-Specific Windows
Annotation Windows
Other Windows

Message-Specific Windows

2 minutes to read

This section describes **Tool Windows** that are *message-specific* and *redockable* windows, and which provide additional details for selected messages. You can display any of the **Tool Windows** indicated below by selecting them in the **Windows** submenu of the global Message Analyzer **Tools** menu, if they are not already displayed.

The following message-specific **Tool Windows** are described in this section:

Message Details Tool Window
Message Data Tool Window
Field Data Tool Window
Message Stack Tool Window
Selection Tool Window
Compare Fields Tool Window

Message Details Tool Window

16 minutes to read

The message **Details Tool Window** is one of the primary windows that is driven by message selection in viewers such as the **Analysis Grid**, **Pattern Match**, **Gantt**, and others. For example, whenever you select a message row in the **Analysis Grid** viewer or in the **Matched Instances** section of the **Pattern Match** viewer, the grid configuration of the **Details** window below the **Analysis Grid** viewer is populated with field names, values, types, and other data associated with the particular message that you selected. In addition, **Summary** information for the selected message displays in the tray of the **Details** window for convenience and improved visibility. You can also hover over the tray to display a tooltip with the **Summary** information, which is handy when the **Summary** line is long.

Viewing Message Details Inline

You can also view message details inline by clicking the stacked-cube icons to the left of any **MessageNumber** designator in the **Analysis Grid** viewer. This can be useful if you want to compare two sets of details. For example, while the details of a selected message row display in the **Details** window, you can click the stacked-cube icon of any other message in the **Analysis Grid** viewer to display its details inline.

> **NOTE**
>
> A *green* stacked-cube icon next to the **MessageNumber** designator in the **Analysis Grid** viewer represents a top-level message, while a *blue* stacked-cube icon represents an Operation, as defined in OPN, which typically represents a grouping of request and response message pairs.

When you click one of the stacked-cube icons, a drop-down displays that has up to four tabs, with each tab containing different information as follows:

- **Fields** — reproduces the grid configuration of field data that appears in the **Details** window below the **Analysis Grid** viewer, as later described.
- **Stack** — displays the origins tree that contains the protocol/module stack, with the messages and fragments that were part of a top-level message or operation. This tab makes it simple to quickly obtain a discrete view of the network layers for any selected top-level transaction.
- **Diagnosis** — displays parsing diagnostic information that includes a diagnosis **Type**, diagnosis **Level**, and a descriptive **Message** for any parsing errors that occurred in a set of trace results. Note that you can also find this information in the **Message Stack Tool Window**.
- **Embedded** — displays when a module definition reflects that a module is embedded; for example, SOAP headers are embedded.

The grid configuration for inline message details that are displayed in the **Fields** tab consist of the following columns of information:

- **Name** — describes the names of the fields contained in the message you selected in the **Analysis Grid** viewer.
- **Value** — specifies the value of each field name that is contained in the message you selected in the **Analysis Grid** viewer.

> **NOTE**
>
> The inline **Details** now include a Bit Field display for **Flag** fields. For example, if you expand the **Flags** field in a TCP message, you will see a visual representation of each flag value in an 8-bit field that appears in parentheses in the **Value** field, for example: (. . .0. . . .) for ACK and (. . . . . .1.) for SYN. Note that the same visual representation appears in the **Details Tool Window**.

- **Type** — specifies the data type for the fields contained in the selected message.
- **Bit Offset** — specifies the offset value in bits for each data field, as measured from the beginning of message data to the beginning of a particular data field.
- **Bit Length** — specifies the value in bits for the length of each data field.

> **NOTE**
>
> The *inline* message details do not interact with the **Message Data Tool Window**. Only if you select a particular field in the **Details** window will you see a corresponding hexadecimal value highlighted in the **Message Data** window.

Using the Details Tool Window Features

For any message row that you select in the **Analysis Grid** viewer, in the **Grouping** viewer when in **Selection Mode**, in the **Matched Instances** section of the **Pattern Match** viewer, or in most other data viewers where you select messages, the message **Details** window displays the grid configuration of data fields that were parsed for that message. In some cases, message selection represents multiple messages. When this occurs, Message Analyzer shows the main details data for the in-focus message. If you previously selected a field in the **Details** window for a particular message type, Message Analyzer will reselect the same field whenever you select another message of the same type. In addition, whenever a field in the **Details** window is selected, the hexadecimal value (**Bit Length**) of that field is highlighted in the **Message Data** window; the resulting display also provides an indication of the relative location (**Bit Offset**) of the field within the packet. As you manually select other data fields in the **Details** window, the hexadecimal value for each selected field is highlighted within the **Message Data** window. These features enable you to obtain a quick view of message details and hexadecimal data as you scroll through a message collection.

You can also filter any column of data in the **Details** window to isolate specific values by applying a **Column Filter**, similar to the way you do this in the **Analysis Grid** viewer, as described in Filtering Column Data. To apply a **Column Filter**, you must click the **Show Column Filter Row** icon just below the **Details** window toolbar and then enter text in one of the columns to filter the data. Thereafter, only the fields that contain the text value that you specified in a particular column will display data. However, in some cases, a search result value for a field will display only if the tree containing the field is open. The column filtering feature provides a convenient way to group data based on field names, values, types, and so on.

Using the Context Menu Commands

You can also create Filter Expressions, add **Analysis Grid** viewer columns, and perform **Group** operations based on data fields in the **Details** window. In addition, you can display the OPN definition for most data fields that you select. Commands for these operations and others are located on a context menu that displays when you right-click any row of field data in the **Details** window. These commands and the actions that result when you select them are described in the list that follows. In this list, the single-quoted *fieldname* values are placeholders for the actual field names that display in the **Name** column of the **Details** window:

- **Add '<fieldName>' to Filter** — captures the name of the selected field within the associated message hierarchy, together with any value that applies, and constructs a Filter Expression that displays in a **Filter** panel text box on the Filtering Toolbar. This action does not apply the filter, as you must manually do that

by clicking the **Apply** button, on a **Filter** panel of the Filtering Toolbar. This action filters the data displayed in any session viewer tab that has focus, for example, an **Analysis Grid** or **Gantt** viewer instance.

- **Add '<fieldName>' as Column** — captures the name of the selected field within the associated message hierarchy and adds a new named column to the current **Analysis Grid** viewer column layout, based on the selected field name. Message Analyzer populates that new data column with information that derives from the current message collection.

> **NOTE**
> An **Analysis Grid** viewer session tab must have the active focus for this command to appear in the **Details** window context menu.

- **Add '<fieldName>' as Grouping** — captures the name of the selected field within the associated message hierarchy and performs a data grouping operation based on the selected field name.

> **NOTE**
> When an **Analysis Grid** viewer session tab is in focus, this command creates a Grouped view of data based on the selected field name. If a Grouped view already exists in the **Analysis Grid** viewer, the new Group that you add will be configured at the lowest level of the nested group configuration. If the single-instance **Grouping** viewer has focus, this command causes a new Group to be added to the current **Grouping** viewer **Layout** at the lowest level of the nested group configuration.

- **Go To '<fieldName>' Definition** — captures the name of the selected field within the associated message hierarchy and causes Message Analyzer to display the **OPN Viewer**. The selected field is highlighted in yellow within the OPN definition code that displays.
- **Include Hex for Numerical Values** — when selected, this option displays hexadecimal values in parentheses next to the numerical values in the **Value** column of the **Details** window. Unselecting this option removes the hexadecimal values.
- **Display Binary Values as** — provides three options that determine the default format in which Message Analyzer displays binary values, for example, in the **Details** window and the **Analysis Grid** viewer. Options include **ASCII**, **Hex**, and **Decimal**.

  You can set one of these as the default value in the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu; however, you can also override the default setting from the **Display Binary Values as** context menu. Message Analyzer installs with **ASCII** as the default setting, which can be useful in exposing additional information, for example, in **Payload** fields that otherwise display values as array elements only when the **Hex** or **Decimal** option is active.
- **Track '<entityName>'** — enables you to track the value of a particular field or property in the **Details** window, depending on which of the following **Details** toolbar buttons is selected. Note that you can hover over the **Details** toolbar buttons with your mouse to display a tooltip that identifies them:
  - **Show all fields for the selected message** — displays all message field **Names** and **Values** in the **Details** window. From here, you can right-click any field and select the **Track '<entityName'** context menu command to cause Message Analyzer to track the field value across multiple message selections.
  - **Show all properties for the selected message**— displays all message property **Names** and **Values** in the **Details** window. From here, you can right-click any property and select the **Track '<entityName'** context menu command to cause Message Analyzer to track the property value across multiple message selections. The properties that appear after you click the properties button

on the **Details** window toolbar exist in two categories:

- **Global** — this category contains annotations and properties that exist under the **Global Annotations** and **Global Properties** nodes in the **Field Chooser Tool Window**. These properties apply to all messages.
- **Message** — this category is named according to a selected message type, for example, SMB, TCP, HTTP, and so on. It contains properties that are specific to the currently selected message only.

Note that the **Global** properties are additional resources that you can use when creating filters, column configurations, and grouping configurations through the context menu that displays when you right-click a property.

Comparing Field Values
Any fields and properties that you have set for tracking values appear in the **Details** window when you click the **Show tracked fields and properties for the selected message** button. You can set a field or property for tracking in the previously specified manner. Once Message Analyzer is tracking a field or property, click the **Details** toolbar button with the footprint icon to display the value/s. Thereafter, you can select different messages in the **Analysis Grid** viewer and observe how values vary across your trace results, for enhanced analysis.

> **TIP**
> If you have the Selection Tool Window set for tracking messages, you can conveniently backtrack to previous message selections to view tracked field and property values.

- **Copy Selected Rows** — enables you to copy one or more selected rows of data in the **Details** window to the clipboard.
- **Copy '<fieldname>'** — enables you to copy a specific field name to the clipboard.

> **TIP**
> You can also use the keyboard shortcuts **Ctrl+C** and **Ctrl+Alt+C** to initiate the **Copy Selected Rows** and **Copy '<fieldname>'** commands, respectively.

Using the Details Toolbar Functions

The toolbar in the upper part of the **Details** window has several functions that provide interactive capabilities, as follows:

- **Pin selection icon** — until this function is activated by clicking its icon, the **Details** window will snap to **Analysis Grid** viewer message selections. This is the default state, as indicated by the hover-over text that reads "**Selection changes are actively tracked**". Otherwise, after clicking the pin selection icon, the current **Details** window is frozen to the values of the currently selected message, as indicated by the hover-over text "**Selection is pinned, selection changes are ignored**". Also, the background of the pin selection icon turns to an amber color in the pinned state.

  In the pinned state, additional message selection in the **Analysis Grid** viewer no longer drives the display of message field data in the **Details** window. This enables you to compare field values with other instances of the **Details 1** window for selected messages, for example, **Details 2**, **Details 3**, and so on.

- **Show all fields for the selected message** — click this icon to show all fields for a message that you selected.

- **Show all properties for the selected message** — click this icon to show all properties for a message that you selected.
- **Show tracked fields and properties for the selected message** — displays a list of fields and/or properties that you have set for tracking, as previously described.
- **Filter funnel icon** — after selecting a **Details** window field, click this icon to create simple view **Filter** code based on the selected field and place it in the Filter Expression text box on the default Filter panel of the Filtering toolbar. Note that the filtering will apply to the in-focus viewer only, for example the **Grouping** viewer or the **Analysis Grid** viewer.
- **Add selected item to Analysis Grid column** — after selecting a **Details** window field, enables you to add a column to the **Analysis Grid** viewer to reflect the values of the selected field for the messages displayed in the grid. Applies to the **Analysis Grid** only, which must be in focus for this command to be enabled on the toolbar.
- **Add selected item as Grouping** — after selecting a **Details** window field or property, enables you to add a new Group to the current **Grouping** viewer **Layout** based on the selected field, providing that the **Grouping** viewer tab has the active focus. Otherwise, if an **Analysis Grid** viewer session tab is in focus, this command creates a grouped configuration of messages in the **Analysis Grid** viewer, based on the selected **Details** window field or property.
- **Search window** — enables you to search for information in the **Details** window based on a field or property name, value, or other text. Filters all data from display except data in the **Details** window grid that matches the search criteria.

Changing Analysis Perspectives with Details Operations

As described earlier, when a particular viewer is in focus, you can right-click any **Details** window field or property and select an appropriate context menu command (see Using the Details Tool Window Features), to change your analysis perspective. You can also use corresponding toolbar buttons to achieve the same result in the following scenarios:

Grouping Viewer in Focus
If this viewer is in focus, you can use the following context menu commands to achieve the indicated results, based on a selected field or property:

- **Add '<fieldName>' to Filter** — enables you to create and add simple view **Filter** code to the Filter Expression text box of the currently displayed **Filter** panel on the Filtering toolbar. You can then apply filtering to the messages that are displayed in the **Grouping** viewer. This is not only a quick way to create a filter, but a convenient method for applying a filter that targets specific data in a set of trace results, to produce a focused message set in a grouped configuration that can expose the cause of a problem.
- **Add '<fieldName>' as Grouping** — enables you to add a new Group to the **Grouping** viewer to enhance the grouped data configuration for better analysis capabilities.

Analysis Grid Viewer in Focus
If this viewer is in focus, you can use the following context menu commands to achieve the indicated results, based on a selected field or property:

- **Add '<fieldName>' to Filter** — enables you to create and add view **Filter** code to the Filter Expression text box on the default Filter panel of the Filtering toolbar. You can then apply filtering to messages that are displayed in the **Analysis Grid** viewer to achieve a focused set of messages that makes it easier to identify problems, by removing messages that do not pass the filtering criteria.
- **Add '<fieldName>' as Column** — enables you to add a new column to the **Analysis Grid** viewer to expose the data for a selected field or property across a set of trace results. You might then perform a

Group command for the new data column to further enhance your analysis perspective.

- **Add '<fieldName>' as Grouping** — enables you to configure nested Group configurations in the **Analysis Grid** viewer to achieve an organized and hierarchical message display that exposes different but related data in the nested groups.

Displaying and Re-docking the Details Tool Window

If the **Details** window is no longer displaying in an Analysis Session, you can redisplay it by selecting the **Details** item from the **Windows** submenu that is accessible from the global Message Analyzer **Tools** menu. Note that you can also undock and reposition the **Details** window by taking advantage of the docking navigation control that displays after you drag the **Details** window by its tab away from its default docking location. You might do this to place the **Details** window adjacent to a particular **Tool Window** or data viewer with which you are working, to create an enhanced view of your data.

Analyzing Data with Multiple Details Tool Window Instances

The **Details** window provides a multi-instance capability that enables you to display up to a maximum of four numerically numbered windows of this type. This feature supports your data analysis process because it enables you to compare different message details data across multiple instances of this window type.

For example, selecting a message in the **Analysis Grid** viewer or **Message Stack Tool Window** can drive the display of details in the **Details 1** window, if it is currently open. Before selecting another related message of interest in the **Analysis Grid** viewer, you can pin or *freeze* the data in the **Details 1** window with the previously described pin selection icon and then open the **Details 2** window from the **Windows** submenu of the global **Tools** menu. You can then select another related message in the **Analysis Grid** viewer or **Message Stack Tool Window** and the message details of this newly selected message can display in the **Details 2** window. You can now easily compare data between the two **Details** window instances, given that the latter **Details 2** window displays immediately to the right of the former **Details 1** instance in this example. You can display up to four **Details** windows to compare the values of four messages that you select in the **Analysis Grid** viewer.

Example Usage Scenario
As a usage example, consider that you have two traces, one from a computer that is performing well and another that is not. You might also have a common message that you sent to both, such as an SMB FileCreate message. In the case of the poorly performing computer, you notice that the FileCreate message seems to be the point where the traffic diverges into slower performance. You could pin the **Details** for this message and then compare them to the computer that is performing well. For example, you might notice that the server response message on the well performing computer has the OpLocks flag set while the poorly performing computer does not. This might mean that the server did not grant an opportunistic lock (OpLock) to the poorly performing computer, which in turn, could result in poor performance.

See Also

Filtering Live Trace Session Results
Writing Filter Expressions
Using the Field Chooser
Using the Analysis Grid Group Feature
Message Data Tool Window

Message Data Tool Window

7 minutes to read

The **Analysis Grid** viewer works in tandem with the message **Details**, **Message Data**, and **Field Data Tool Windows**. If you select any message row in the **Analysis Grid** viewer, message data can be displayed in the following ways:

- The hexadecimal value of a message field may be highlighted in the **Message Data** window.
- The name of a message field may be selected in the **Details** window grid; the grid also contains the value, type, bit offset, and bit length of each field in a message.
- The decimal or parenthetical hexadecimal value of any selected message field in **Details** can display in the **Field Data** window.

Tool Window Interaction

If you select a particular field in the **Details** window, Message Analyzer highlights the hexadecimal value of that field in the **Message Data** window. The **Details** window also works interactively with **Field Data** window. For example, if you select a field in the **Details** window, the **Field Data** tab displays the decimal and hexadecimal values of the selected field. These capabilities enable you to quickly assess the value of any field contained in any message.

Using the Controls and Commands

The **Message Data** window is accessible as a selectable tab below the **Analysis Grid** viewer, as is the **Field Data** window. The **Message Data** window has several right-click context menu commands that provide options for copying, displaying, and saving hexadecimal, ASCII, and binary data, as described below.

- **Copy** — enables you to copy several forms of data in the **Message Data** window, as follows:
  - **Hex** — a command that enables you to copy any selected Hex value in the **Message Data** window.
  - **ASCII** — a command that enables you to copy the ASCII values of selected hexadecimal values in the **Message Data** window.
  - **Hex & ASCII** — a command that enables you to copy ASCII and hexadecimal values for a selected block of hexadecimal values in the **Message Data** window.
- **Save Selected Bytes As...** — enables you to save selected trace data in binary format (*.bin), whether the Message Data window is configured to display data in Hex, ASCII, or Binary format. This feature can work together with the **Select All** command in this context menu so that you can save all the data in binary.
- **Display Options** — allows you to display any combination of hexadecimal, ASCII, and binary values in the **Message Data** window, as follows. Note that one value is always selected:
  - **Hex** — selection and slider controls that enable you to display hexadecimal data values of a message or message field in variable font sizes.
  - **ASCII** — selection and slider controls that enable you to display ASCII data values of a message or message field in variable font sizes.
  - **Binary** — selection and slider controls that enable you to display binary data values of a message or message field in variable font sizes.

- **Column** — a drop-down control that enables you to specify the column width of the **Message Data** window in terms of the number of bytes displayed. The default setting is 16 bytes.

- **Select All** — click this command to highlight all the data displayed in the **Message Data** window. This feature is designed to work with the **Save Selected Bytes As...** command to save all selected data in binary format.

> **NOTE**
> The **Field Data** window also has several right-click menu commands that are described in the Field Data Tool Window topic.

The **Message Data** window has several labels in the tray area of the window that display the following information:

- **Byte Count** — provides an indication of the number of bytes of data that are currently selected.
- **Message Offset** — indicates how far the starting point of field data or a highlighted block of hexadecimal values is offset, in bytes, from the beginning of a message.
- **Protocol Offset** — indicates how far the starting point of field data or a highlighted block of hexadecimal values is offset, in bytes, from the beginning of a protocol message, which starts at the point where the stripped header (un-highlighted values) of a previous lower layer ends.

Viewing Field and Message Data

You can view the fields and values of any top-level parent message node that you select in the **Analysis Grid** viewer, and you can also view the fields and values of any selected child message node that is part of the origins tree under its parent node. To expose this data, expand the parent message node and each child message node in succession to display the underlying stack layers that make up the origins tree. For example, you might have a TCP message as a parent node in a **Local Network Interfaces** trace, with subsequent child nodes for the **IPv4**, **Ethernet**, **NdisProvider**, and **Etw** messages that make up the origins tree. If you select a message row in the **Analysis Grid** viewer for any one of these messages and you have not yet selected a field in the **Details** window, all the data contained in the selected message is highlighted in the **Message Data** window. Moreover, if you select a particular field of the message in the **Details** window, Message Analyzer highlights the hexadecimal value of the field in the **Message Data** window. At the same time, the value of the field displays in the **Field Data** window.

You can also examine the headers and payloads of any particular message layer and you can even observe how the process of de-encapsulation is represented, as you upwardly traverse the message stack. For instance, using the previously indicated TCP example, perform the following steps to visualize how de-encapsulation took place:

To visualize message layer de-encapsulation

1. Perform a **Local Network Interfaces** trace and let Message Analyzer run long enough to capture some top-level TCP messages.
2. Expand a parent TCP operation node and then expand the message nodes in the origins tree to expose the message stack.
3. Select the **Etw** node in the **Analysis Grid** viewer and then select the **Payload** field in the **Details** window.

   The hexadecimal values for the **Payload** field are highlighted in the **Message Data** window.
4. Select the **NdisProvider** node in the **Analysis Grid** viewer and then select the **FrameData** field in the **Details** window.

   The hexadecimal values for the **FrameData** field are highlighted in the **Message Data** window.

5. Select the **Ethernet** node and then select the **MacClientData** field in the **Details** window.

   The hexadecimal values for the **MacClientData** field are highlighted in the **Message Data** window.

6. Select the **IPv4** node and then select the **Payload** field in the **Details** window.

   The hexadecimal values for the **Payload** field are highlighted in the **Message Data** window.

7. Select the **TCP** node and then select the **Payload** field in the **Details** window.

   Only the TCP header data (the IPv4 payload) remains, with no hexadecimal data highlighted, while the **Payload** field is selected. To see this more clearly, go back and select the **Etw** message in the **Analysis Grid** viewer, and then in succession, select the **NdisProvider**, **Ethernet**, **IPv4**, and **TCP** messages. As you do, you can observe the header data being stripped out at each layer, which reflects the PEF Runtime decoding process that de-encapsulates captured messages, as illustrated in the figure that follows.

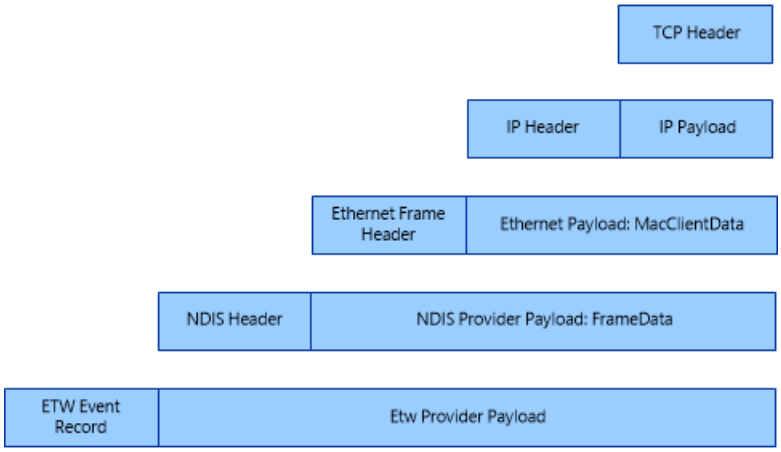

Figure 66: Message Analyzer Layer De-encapsulation example

Displaying the Message Data Tool Window

If the **Message Data** window is no longer displaying in an Analysis Session, you can redisplay it by selecting the **Message Data** item from the **Windows** submenu, which is accessible from the global Message Analyzer **Tools** menu. You can also undock and reposition the **Message Data** window by taking advantage of the docking navigation control that displays after you drag the **Message Data** window away from its default docking location by its tab. You might do this to move the **Message Data** window adjacent to another **Tool Window** or data viewer with which you are working, to enhance your data analysis perspectives.

Analyzing Data with Multiple Message Data Tool Window Instances

The **Message Data** window provides a multi-instance capability that enables you to display up to a maximum of four numerically numbered windows of this type. This feature supports your data analysis process because it enables you to compare the hexadecimal data of different messages across multiple instances of this window type.

For example, selecting a message in the **Message Stack 1** window, **Analysis Grid** viewer, **Grouping** viewer, or other data viewer, can drive the display of hexadecimal values in the **Message Data 1** window, if it is currently open. Before selecting another related message of interest in a particular viewer, you can pin or *freeze* the data in the **Message Data 1** window by clicking the pinning icon on its toolbar and then open the **Message Data 2** window from the **Windows** submenu that is accessible from the global **Tools** menu. You can then select another related message in a particular viewer. Thereafter, the hexadecimal values of any message that is automatically selected in the **Message Stack** window or of any field that is automatically selected in the **Details** window are displayed in the **Message Data 2** window. You can now easily compare data between the two **Message Data** window instances. You can display up to four **Message Data** windows to compare the values of four messages

that you select in a particular viewer.

See Also

Message Details Tool Window

Field Data Tool Window

3 minutes to read

The **Field Data Tool Window** is a field-specific and dockable window that is interactively driven by field selection in the **Details Tool Window**, which is in turn driven by message selection in the **Analysis Grid** and other data viewers. Interaction between the **Analysis Grid** viewer and the **Details** and **Field Data** windows occurs only within the context of individual sessions, rather than across different sessions.

Displaying Field Data

In most cases, the **Field Data** window reproduces the value of a selected field that is specified in the **Value** column of the **Details** window, which could be a string, numerical, binary, or other value. For example, if you select a field in the **Details** window for a TCP message that is selected in the **Analysis Grid** viewer, the value of that field also displays in the **Field Data** window. In cases where the information is reproduced, the **Field Data** window will indicate that there is "*No alternate presentation available*" for these fields.

However, there are some cases where more complex field information can display in the **Field Data** window. This can enable you to very quickly assess the type of payload data a message is carrying. For instance, if you select a TCP message in the **Analysis Grid** viewer and the **Payload** field in the **Details** window, the **Field Data** window usually displays information in text/plain format. If the TCP payload is an HTTP message, then the **Field Data** window might display HTML fragments. Moreover, if you select the **Payload** field in the **Details** window for an HTTP message selected in the **Analysis Grid** viewer, the **Field Data** window might display formats that include the following: text/html, text/css, font/eot, application/x-javascript, text/javascript, and application/json; or different media formats such as image/gif, image/jpeg, image/png, image/bmp, and so on. Also, XML is another format that can display in the **Field Data** window.

> **NOTE**
> The **Field Data** window indicates the data format or presentation type by displaying a label value in the lower-left corner of the window.

Decoding URLs

The **Field Data** window now provides decoding of URLs by removing the special encoding characters from any URL string. For example, when you select the **Uri** field in the **Details** window for an HTTP message, the **Field Data** window displays an upper section named **Actual Value** and a lower section named **Decoded Value**. The latter section shows the decoded value of the **Uri** with the encoding characters removed, while the original value of the **Uri** is retained in the **Actual Value** section. A typical result might be that all **Uri** spacing characters such as "%20" are removed.

Correlating Field Details

You can also correlate the hexadecimal details for most fields selected in the **Details** window with the values displayed in the **Field Data** window, by observing them in the **Message Data Tool Window**. For example, to view the hexadecimal details for a TCP **Payload** field value that is displaying in the **Field Data** window, click the **Message Data** tab.

Working With the Field Data Tool Window

If the **Field Data** window is no longer displaying, you can redisplay it by selecting the **Field Data** item in the **Windows** submenu that is accessible from the global Message Analyzer **Tools** menu. You can also undock and reposition the **Field Data** window by taking advantage of the docking navigation control that displays after you drag the **Field Data** window by its tab away from its default docking location. You might do this to move the **Field Data** window adjacent to another **Tool Window** or data viewer with which you are working, to enhance your data analysis perspectives.

After the **Field Data** window is displayed, you can use it in the previously specified manner. When data is displaying in this window, you can select the following right-click context menu items to perform the indicated tasks:

- **Save Text As ...** — Message Analyzer interprets the type of field data displaying in the **Field Data** window and causes the **Save As** dialog to open with a file extension specification that correlates with the interpreted data type. For example, if the displaying field data is **text/html**, the **Save As** dialog opens with the **Save as type** drop-down selection set to **HTML Files (*.html)**.
- **Open in Default Application** — Message Analyzer interprets the type of field data displaying in the **Field Data** window and causes a default application to open for the associated type of field data. For example, if the displaying field data is indicated as **text/css**, the default HTML editor on your system opens, such as SharePoint Designer or Expression Web, to enable you to save the data in that application.
- **Copy Text** — enables you to simply copy the text to the Clipboard for further processing.

Message Stack Tool Window

5 minutes to read

The **Message Stack Tool Window** is an interactive, message-specific, and dockable window where the display of details is driven by message selection in a viewer such as the **Analysis Grid**, **Grouping**, or other viewer. The result is an independent display of the full stack for any message selected in a particular viewer. This enables you to quickly observe and explore the network architecture of any message, without having to manually open different parts of the stack (origins messages) in the **Analysis Grid** viewer. The **Message Stack** window contains a toolbar that provides different buttons for changing the format of the message stack view, as follows:

- **Visual stack** — the default view which provides data that is organized into a simple visual stack display along with summary information.
- **Table** — provides a table with data in three columns consisting of the top-level **Message** number, **Module** name, and a **Summary** description for the selected message.
- **Tree with field information** — provides similar data with a different organization of summary information.
- **Origins** — the toolbar in the **Message Stack** window also provides a label that indicates the **Origins** count, which is a measure of the number of origins trees under a top-level message node.

  For example, there are typically two origins trees under an operation, one that represents a request stack and one that represents the response stack. Under other top-level messages, there is typically one origins tree. However, if there are message fragments, the **Origins** count can be higher.

Analyzing Stack Data

The **Message Stack** window can act as a central hub from where you can quickly correlate additional message details and contexts. This capability enables you to enhance your data analysis perspectives because you can do the following:

- Drive the data displayed in the **Details Tool Window**. For example, by selecting a network/message node in the **Message Stack** window, the **Details** window is automatically populated with the field data of the selected message.
- Drive selection of top-level messages in the **Analysis Grid** viewer, which can show the context of the surrounding message configuration in the **Analysis Grid** viewer.
- Drive message selection in the **Analysis Grid** viewer of any origins message under a top-level parent message node, providing that the origins tree is in the expanded state for the particular top-level message. Otherwise, no messages will be highlighted when you select any origins message in the **Message Stack** window.
- Compare **Message Stack** information for one message in the **Analysis Grid** viewer with the message stack of a second message in the **Analysis Grid**, by clicking the cubed icon of the second message and comparing the data on the **Stack** tab of the inline message details in the **Analysis Grid** viewer with the data displaying in the **Message Stack** window.

Using the Message Stack Window Interactive Features

After you select a top-level message row in the **Analysis Grid** viewer, the entire message stack for the selected message displays in the **Message Stack** window. If you expand all the message nodes underneath the top-level

message in the **Analysis Grid** viewer and then select any of its child messages (origins tree messages), the **Message Stack** window highlights those child messages within the stack in light blue. Likewise, if you select any message in the **Message Stack** window, the corresponding message rows in the **Analysis Grid** viewer are automatically highlighted, provided that the nodes for those messages are in the expanded state in the **Analysis Grid** viewer; otherwise, no messages will be highlighted. The selection in the **Message Stack** window also displays field data in the **Details** window grid for the selected node in the stack, even if the node of the associated top-level message in the **Analysis Grid** viewer is unexpanded.

Using the Message Stack Window Context Menu Commands

The **Message Stack** window has several context menu commands that display only when you right-click a message row in the Table view. These commands and the actions they invoke are described as follows:

- **Show Toolbar** — alternately displays and hides the **Message Stack** window toolbar.
- **Copy Selected Rows** — enables you to copy the data for one or more selected message rows in the **Message Stack** window. You can also use the keyboard shortcut `Ctrl+C` to execute this command.
- **Copy '<columnName>'** — enables you to copy the data in any column that you right-click in the **Message Stack** window. You can also use the keyboard shortcut `Ctrl+Alt+C` to execute this command.

Displaying the Message Stack Tool Window

If the **Message Stack** window is no longer displaying in an Analysis Session, you can redisplay it by selecting the **Message Stack** item in the **Windows** submenu, which is accessible from the global Message Analyzer **Tools** menu. You can also undock and reposition the **Message Stack** window by taking advantage of the docking navigation control that displays after you drag the **Message Stack** window by its tab away from its default docking location. You might do this to move the **Message Stack** window adjacent to another viewer or **Tool Window** with which you are working, to enhance your data analysis perspective.

Analyzing Data with Multiple Message Stack Tool Window Instances

The **Message Stack** window provides a multi-instance capability that enables you to display up to a maximum of four windows of this type. This feature supports your data analysis process because it enables you to compare the top-level and origins data of different messages across multiple instances of this window type.

For example, selecting a message in the **Analysis Grid** viewer drives the selection of the corresponding message in the **Message Stack 1** window, if it is currently open. Before selecting another related message of interest in the **Analysis Grid** viewer, you can pin or *freeze* the data in the **Message Stack 1** window and then open the **Message Stack 2** window from the **Windows** submenu in the global **Tools** menu. You can then select another related message in the **Analysis Grid** viewer and the newly selected message is automatically selected in **Message Stack 2** window. You can now easily compare data between the two **Message Stack** window instances.

Selection Tool Window

7 minutes to read

The **Selection Tool Window** is an interactive and dockable window that enables you to readily expose message selection in several Message Analyzer data viewers, which is especially advantageous in cases where message selection could be ambiguous. For example, in a data viewer such as the **Gantt** chart, message selection is not as obvious as other viewers such as the **Analysis Grid** or **Interaction** chart. The **Selection** window makes your selection more obvious because it exposes message selection in its window as a table of information that appears separate from the viewers in which you make selections. Integral to this capability and equally as important is the ability of the **Selection** window to maintain the context of multiple message selection in the following ways:

- Within a specific viewer, such as the **Analysis Grid**.
- Across multiple viewers or multiple instances of the same viewer type, in the same session.
- Across multiple viewers in different sessions.

The **Selection** window provides a separate space that independently monitors and displays the selection of messages in multiple viewers and keeps track of such selections, so that you can thereafter back- and forward-navigate among messages as necessary, while maintaining the context of previously selected messages. This enables you to undo any message selection that you made accidentally, without affecting other messages that you selected. The **Selection** window also provides various **Column Layouts** that can expose different message fields. All of these features enable you to improve message correlation and analysis capabilities.

Modes of Operation

The **Selection** window provides the following modes of operation that diversify the scope of message selection:

- **Single viewer selection mode** — this mode is the default state in which the **Selection** window initially displays. In this mode, both the **Activate Session Viewer** and **Navigate Across Sessions** buttons are disabled, as indicated by a subtle blue background color that displays when you hover over them with your mouse. The name of each button also displays in a popup when you hover over them. In this mode, you can only navigate among previously selected messages in the currently active data viewer. For example, as you select multiple messages in a single data viewer such as the **Analysis Grid**, the **Selection** window builds a collection of the selected messages. Thereafter, you can use the **Go back** and **Go forward** arrow-buttons on the **Selection** window toolbar to navigate through the messages that you selected in the *currently active data viewer* only.
- **Activate session viewers selection mode** — this mode is enabled by clicking the **Activate Session Viewer** button on the toolbar of the **Selection** window. When this mode is active, the **Activate Session Viewer** button displays with an amber background color. In this mode, the **Selection** window tracks all messages that you select in a single session by building an associated message collection, where the data is displayed in one or more data viewers — for example, one or more **Analysis Grid** and **Gantt** viewer instances in the same session. After you select a number of messages, you can use the **Go back** and **Go forward** arrow-buttons on the **Selection** window toolbar to navigate through the message collection that is displayed in *current session viewers only*. To disable this mode, click **Activate Session Viewer** button again to remove the amber background color.

  As a usage example for this mode, you might have multiple book marked messages that you have selected in the **Bookmarks Tool Window**, where such selection in turn drives message selection in the **Analysis Grid** viewer. Furthermore, because the **Selection** window interacts with the **Analysis Grid**, it responds to messages that are selected in the grid—and also to the manner in which they are selected, as in single or

multiple message selection—by adding those messages to the **Selection** window. While reviewing a group of book marked and selected messages in the **Analysis Grid** viewer that are distributed throughout a trace, you can lose the multiple selection context in the grid if you select any single message in the group for further inspection. With the **Selection** window, you can select individual messages without resetting the selection context in the **Analysis Grid** viewer. Moreover, as you select messages in the **Selection** window, it drives message selection in the **Analysis Grid** viewer by placing the selection arrow to the left of each selected message; however, the original multiple message selection context in the grid is not lost as you do this.

- **Navigate across session viewers selection mode** — this mode is enabled by clicking both the **Navigate Across Sessions** and the **Activate Session Viewer** buttons on the toolbar of the **Selection** window. When this mode is active, both of these buttons display with an amber background color. In this mode, the **Selection** window tracks all messages that you select in multiple sessions by building an associated message selection collection, where the data is displayed in one or more viewers in each session. After you select a number of messages in different viewers from different sessions, you can use the **Go back** and **Go forward** arrow-buttons on the **Selection** window toolbar to navigate through the message collection that spans across *all session viewers* in which you selected messages. To disable this mode, click these same buttons again to remove the amber background color.

Example Use Cases

As a simple usage example when navigating across sessions, you might have multiple log files or other saved traces that contain related data in separate **Analysis Grid** viewer session tabs that you need to analyze. You can then use the **Go To Message** function on the **Analysis Grid** toolbar to jump to specific messages in each session tab that you know have related data. Each time you jump to a message, it is displayed in the **Selection** window, which in turn keeps track of all your selections, as described earlier. You can then use the **Go back** and **Go forward** arrow-buttons on the toolbar of the **Selection** window to navigate back and forth between the messages with related data for a comparative analysis. For each selected message that you navigate to, you can view the associated data in the **Message Stack**, **Details**, and **Message Data Tool Windows** to correlate the details you need to analyze. To further facilitate this scenario, you can drag the **Selection** window tab away from its default docking location in the lower right corner of the Message Analyzer UI so you can more easily expose the **Message Data** window in its default docking location.

Another usage example that applies to any mode is if you click a message in the **Analysis Grid** viewer, possibly to bring it into focus, and you inadvertently undo a previous selection that is important to the analysis in which you are engaged. With the **Go back** or **Go forward** feature, you can resume navigation among previous message selections, regardless of whether the context of the previous message selection was lost.

Changing Column Layouts in the Selection Tool Window

The **Selection** window also provides a **Default** column layout that has some similarities with the default column layout of the **Analysis Grid** viewer. However, you can change the column layout within the **Selection** window for analysis purposes. To do this, right-click anywhere within the **Selection** window to display the **Column Layout** context menu, which contains all the column layouts that are available from the **Layout** drop-down list on the **Analysis Grid** toolbar; an **Example View Layout** is also provided. By selecting a layout item from the **Column Layout** submenu, the specified layout replaces the current column configuration of the **Selection** window.

This feature enables you to conveniently maintain the context of the current **Analysis Grid** column **Layout**, while simultaneously and independently displaying a different layout in the **Selection** window that is tailored for a more detailed and on-the-spot analysis of specifically selected messages. For example, you could select a TCP message in the **Analysis Grid** viewer and display it in the **Selection** window. Thereafter, you can right-click the message in the **Selection** window and select the **TCP** item from the **Column Layout** context menu to display fields that are important for TCP analysis, such as the following:

- **DiagnosisTypes**
- **Timestamp**
- **TimeDelta**
- **Source address**
- **Destination address**
- **SourcePort**
- **DestinationPort**
- **PayloadLength**
- **SequenceNumber**
- **AcknowledgementNumber**
- **Window**
- **Summary**

> **TIP**
> If you want to create your own view **Layouts** that you can select from the **Selection** window **Column Layout** menu, you will need to create a new **Layout** in the **Analysis Grid** viewer and save it from the **Layout** drop-down list on the **Analysis Grid** toolbar. For more information, see Applying and Managing Analysis Grid Viewer Layouts.

Displaying and Re-docking the Selection Tool Window

If the **Selection** window is not already displayed, you can display it by clicking the **Selection** item in the **Windows** submenu, which is accessible from the global Message Analyzer **Tools** menu. Note that you can redock the **Selection** window by using the docking navigator control that displays after you undock the window from its default location. Undock the **Selection** window by dragging it by its tab away from the default location.

You can use this feature to facilitate a convenient location for your analysis processes. For example, you might undock this window from its default location and reposition it alongside another **Tool Window** or data viewer, or you can allow the window to float in a chosen location outside the Message Analyzer user interface. This is also possible for any other **Tool Window** or data viewer.

Compare Fields Tool Window

2 minutes to read

Message Analyzer enables you to compare the field data of any two messages that you select in the **Analysis Grid** viewer for detailed analysis of field values. This functionality is provided in the **Compare Fields Tool Window**, which is a preview feature that you must enable on the **Features** tab of the **Options** dialog in order to use it. After enabling this preview feature and restarting Message Analyzer, you can access it from the **Windows** submenu of the global Message Analyzer **Tools** menu.

Comparing Message Field Data

To begin, open the **Compare Fields** window from the **Windows** submenu as previously indicated, to display it in the lower right sector of the Message Analyzer user interface. Assuming that you have a set of trace results displaying in the **Analysis Grid** viewer, perform the following steps:

1. In the **Analysis Grid** viewer, select a message containing the fields that you want to compare with those of another message of the same type.
2. In the **Compare Fields** window, right-click anywhere in the window space and select the **Set Current as Baseline** command from the context menu that appears, to set the currently selected message fields and values as a reference in the window.

   Field and **Baseline** columns with values derived from the selected baseline message appear in the **Compare Fields** window.
3. In the **Analysis Grid** viewer, select a second message of the same type which contains the fields that you want to compare in value to the baseline message fields.

   A **Current** column with field values derived from the currently selected message appears in the **Compare Fields** window.
4. Scroll down through the data and note that fields in the **Current** column which contain values that are different from those in the **Baseline** column are highlighted in red, thus providing you with an instant evaluation of differences between message field values that could be critical to analysis.

Session-Specific Windows

2 minutes to read

This section describes interactive **Tool Windows** that are *session-specific* and dockable windows, and which provide additional details for selected sessions. You can display the session-specific **Tool Windows** indicated below by selecting them in the **Windows** submenu that is accessible from the global Message Analyzer **Tools** menu, if they are not already displayed. However, for the **Diagnostics Tool Window** to appear in the **Windows** submenu, you must first enable it as a preview feature on the **Features** tab of the **Options** dialog, which is also accessible from the global Message Analyzer **Tools** menu. You must then restart Message Analyzer.

The following session-specific **Tool Windows** are described in this section:

Diagnostics Tool Window
Decryption Tool Window

Diagnostics Tool Window

5 minutes to read

In troubleshooting scenarios, quick access to diagnostic message information can help you to rapidly identify the root cause of an issue. However, diagnosis messages are typically embedded within individual captured messages, which can make them very difficult to find. Moreover, there can be multiple diagnosis messages within individual captured messages and the same diagnosis message might appear in multiple captured messages. When a diagnosis message is associated with either a top-level parent message or a child (origins tree) message in the **Analysis Grid** viewer, Message Analyzer provides an initial indication of the diagnosis message at the top-level. This means that even when a diagnosis message is associated with a child message, the diagnosis message *indication* is bubbled up to a top-level parent message, which is beneficial to analysis. However, this display configuration still requires you to go through some manual expansion of message nodes to get at the diagnosis message/s of interest.

Exposing Diagnosis Messages

Message Analyzer has several data manipulation features that you can use to expose diagnosis messages, which include **Grouping** and sorting the **DiagnosisTypes** column in the **Analysis Grid** viewer, but there are some disadvantages in using these techniques. When you sort the **DiagnosisTypes** column to bubble up diagnostic messages that occurred in a trace, or **Group** the column to organize diagnosis messages according to their **DiagnosisType**, you can lose the context of the original surrounding messages, which can often call attention to where or why certain issues occurred. As an alternative to these techniques, you can utilize the **Diagnostics Tool Window**, which flips the relationship of captured-to-diagnosis messages by making diagnosis messages the focus. This enables you to view diagnosis messages independent of the **Analysis Grid** viewer without having to expand message nodes to find them, while at the same time driving selection of parent messages in the **Analysis Grid** viewer that contain the diagnosis messages. You can also drive and synchronize the display of parent message details in other **Tool Windows** through selection of diagnosis messages in the **Diagnostics** window, as described ahead.

> **IMPORTANT**
>
> The **Diagnostics** window is a preview feature. If you want to use it, you will need to select it on the **Features** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. You will then need to restart Message Analyzer. Thereafter, the **Diagnostics (Preview)** window will appear in the **Windows** submenu that is accessible from the global **Tools** menu.

Enhancing Diagnostic Capabilities

The **Diagnostics** window is a session-specific and dockable window that simplifies the means of locating and analyzing embedded diagnosis messages that might normally be hard to find. It summarizes the different types of diagnosis messages in a selected session, which includes **Application**, **Validation**, **InsufficientData**, and **Parsing** errors. It also provides other information such as **Module**, **Count**, and diagnosis **Message** text, to assist your analysis process. By making this information quickly accessible and enabling you to associate other message details, the **Diagnostics** window enhances your troubleshooting experience because you can:

- Focus on diagnosis messages only.
- View a summary of diagnosis message group counts.
- Synchronize diagnosis message selection with the top-level parent messages in the **Analysis Grid** viewer.

- Drive the display of data in the following **Tool Windows**:
  - Message details in the **Details Tool Window**.
  - Hexadecimal field or message data in the **Message Data Tool Window**.
  - Message selection in the **Message Stack Tool Window**.

> **NOTE**
> The **Diagnostics** window drives message selection in the **Analysis Grid** viewer, but the **Analysis Grid** does not drive diagnosis message selection in the **Diagnostics** window.

Viewing Diagnostic Data

The **Diagnostics** window contains four sortable columns that display diagnosis message information as follows:

- **Type** — specifies the type of diagnosis message, such as **Application**, **Validation**, **InsufficientData**, or **Parsing**.
- **Module** — specifies the name of the module or protocol for the parent message that contains the diagnosis message.
- **Message** — specifies the descriptive text of the diagnosis message.
- **Count** — specifies the number of parent messages in the **Analysis Grid** viewer that contain the selected diagnosis message.

More Information
To learn more about the meaning of diagnosis messages, see the Diagnosis Category topic.

Using the Context Menu Commands

The **Diagnostics** window also contains a context menu that enables you to execute the following commands to accomplish the indicated tasks:

- **Ignore Selected Modules** — removes all diagnosis messages from the **Diagnostics** window display that match the module of the selected diagnosis message, so you can focus on the diagnosis messages of other modules.
- **Ignore Selected Messages** — removes all selected diagnosis messages from the **Diagnostics** window display so you can focus on specific diagnosis messages.

  > **NOTE**
  > The toolbar of the **Diagnostics** window has an **Ignore** drop-down list that also contains the two previously described command functions. In addition, you can click the **Reset** button on the toolbar to restore all ignored items, as described next.

- **Reset Ignored Messages** — restores all messages or modules that were previously ignored.
- **Open Selected Messages** — displays the top-level parent messages (including origins) for all selected diagnosis messages in a separate **Analysis Grid** viewer tab.

> **NOTE**
>
> The toolbar of the **Diagnostics** window contains an **Open** button that enables you perform the same function as **Open Selected Messages**.

- **Copy Selected Rows** — copies the data for the selected rows to the Clipboard.
- **Copy <columnName>** — copies the data in a selected column for a selected diagnosis message to the Clipboard.

Displaying the Diagnostics Tool Window

You can open the **Diagnostics Tool Window** by selecting the **Diagnostics (Preview)** item from the **Windows** submenu that is accessible from the global Message Analyzer **Tools** menu. As indicated earlier, because the **Diagnostics** window is a preview feature, make sure that you first select it on the **Features** tab of the **Options** dialog and then restart Message Analyzer to make it available.

> **NOTE**
>
> After you open the **Diagnostics** window for analyzing trace results, you might notice that a scrolling, marquee-style progress indicator displays for larger message sets while Message Analyzer iterates through the current session data, retrieves the diagnosis messages, and configures a display of summary data.

Repositioning the Diagnostics Tool Window

As with other **Tool Windows**, you have the option to undock and reposition the **Diagnostics** window by taking advantage of the docking navigation control that displays after you drag the **Diagnostics** window away from its default docking location by its tab. You might do this to configure a location for the window that better suits your purposes, for example, to obtain context by correlating message selection in the **Diagnostics** window with message selection in the **Analysis Grid** viewer, as the former drives the latter.

Decryption Tool Window

2 minutes to read

The Message Analyzer **Decryption** feature enables you to view messages that are encrypted with the Transport Layer Security (TLS) and Secure Sockets Layer (SSL) protocols, for example, messages from the HTTP and Remote Desktop (RDP) protocols. To decrypt messages that were captured on a specific server, Message Analyzer requires a valid certificate and password for such a server. Also, you must provide these to Message Analyzer prior to loading data from saved input files through a Data Retrieval Session or before running a Live Trace Session that you expect to contain messages that you want to target for decryption. Thereafter, Message Analyzer can decrypt those messages by using the server certificate and password that you provided. To view the results of a decryption session, you will use the **Decryption Tool Window**.

Displaying Decryption Data

To display the **Decryption** window, select the **Decryption** item in the **Windows** submenu that is accessible from the global Message Analyzer **Tools** menu. The **Decryption** window is a session-specific, interactive, and dockable window that simplifies the means of locating and analyzing decrypted messages that might typically be hard to find in a large trace. The **Decryption** window is session-specific because the data it displays is driven by session selection; for example, this might mean selecting an **Analysis Grid** viewer session tab, as described in Data Viewer Concepts. Moreover, if you have decrypted data in an **Analysis Grid** viewer session tab and you select it, the **Decryption** window displays summary data for the selected decryption session. If you have multiple sessions that have decrypted data, selecting a viewer tab for any of them causes the **Decryption** window to snap to the selected session data. This enables you to have quick access to comprehensive decryption information for each decrypted conversation in the selected session.

Assessing Decryption Session Results

In the **Decryption** window, a separate row displays for each message that Message Analyzer attempted to decrypt and status information for each message is also provided to enable you to quickly assess the results of a decryption session. Selection of a message row in the **Decryption** window automatically selects that message in the **Analysis Grid** viewer. This also causes field data in the **Details Tool Window**, hexadecimal data in the **Message Data Tool Window**, and message layer data in the **Message Stack Tool Window** to snap-to the selection, provided that such **Tool Windows** are displayed. This interactive display of associated data immediately provides you with the diagnostic and analysis perspectives by which you can accurately assess your decrypted messages and the supporting origins stack.

More Information

To learn more about the decryption session status information that is provided in the **Decryption** window, see Analyzing Decryption Session Data.

See Also

Decrypting TLS and SSL Encrypted Data

Annotation Windows

2 minutes to read

This section provides the details for using the **Bookmarks** and **Comments Tool Windows** to add annotations to one or more messages that you select in the **Analysis Grid** viewer. These windows are described in the following topics of this section:

Bookmarks Tool Window
Comments Tool Window

Bookmarks Tool Window

14 minutes to read

Message Analyzer provides a **Bookmarks Tool Window** that enables you to configure the following types of bookmarks:

- **Message Bookmarks** — enables you to specify a bookmark for any single message or group of messages that you select in the **Analysis Grid** viewer. You typically set such bookmarks in a trace for quick location of messages that have some related importance or critical context for the data analysis process. Bookmarks make it convenient to locate these messages at any time, since you can save them with your message data in the *.matp format. It also makes it convenient to share the trace with others along with the bookmarks that you configured to highlight the important focus areas.
- **Pattern Bookmarks** — enables you to specify bookmarks for one or more matched instance messages in the **Pattern Match** viewer, after executing a Pattern expression against a set of trace results. You can set these types of bookmarks to quickly locate messages for specific matched instances of a Pattern expression, as described in Viewing Matched Instance Message Data, although the method to do so is different than you would do in the **Analysis Grid** viewer, as described in Viewing Bookmarked Data.

Displaying the Bookmarks Tool Window

If the **Bookmarks** window is not displaying in an Analysis Session, you can open it by clicking the **Bookmarks** item in the **Windows** submenu that is accessible from the global Message Analyzer **Tools** menu. You can also undock and reposition the **Bookmarks** window in a more convenient location by taking advantage of the docking navigation control that displays after you drag the **Bookmarks** window by its tab away from the default docking location.

Configuring Bookmarks

After you display the **Bookmarks** window, you can configure bookmarks for messages in the **Analysis Grid** viewer and for matched instances in the **Pattern Match** viewer. To begin, select a message or matched instance as appropriate for the entity that you want to bookmark, and then select the corresponding command in the **Add** drop-down menu on the toolbar of the **Bookmarks** window:

- **Message Bookmark** — creates a bookmark for one or more messages that you select in the **Analysis Grid** viewer. The **Bookmarks** configuration grid is populated with a row of data that represents the selected message/s.
- **Pattern Bookmark** — creates a bookmark for one or more matched instances that you select in the **Pattern Match** viewer. The **Bookmarks** configuration grid is populated with a row of data that represents the selected matched instance/s that contain the book marked messages.

The **Bookmarks** grid contains the following information columns:

- **Name** — this sortable column provides a default name value of "**Selected messages**" or "Pattern Group <n>", where <n> is a number placeholder. You can modify this column with a name of your choice, either by double-clicking the existing name in this text box or by right-clicking the text and selecting the **Edit 'Name'** item in the context menu that appears. The existing text is then highlighted for edit mode.
- **References** — specifies a single message number or a comma-separated list of the message numbers for the messages you selected in the **Analysis Grid** viewer, or displays a **Selected Pattern Match Instance** (containing one or more messages) that you selected in the **Matched Instances** section of the **Pattern**

Match viewer for which you are configuring a bookmark.

- **Category** — this sortable column is empty by default, but you can specify your own **Category** name for any bookmark. You can do this either by double-clicking the **Category** text box for a particular bookmark row in the **Bookmarks** configuration grid, or by right-clicking the **Category** text box in a bookmark row and selecting the **Edit 'Category'** item in the context menu that appears. You can then create a new **Category** name or edit an existing one.
- **Messages** — this column specifies the number of messages that exist in a bookmark row. To view these messages, click on the book icon.
- **Flag** icon — this sortable column is empty by default, but you can configure it with a chosen flag color to indicate a self-defined level of importance, interest, or critical context that a bookmarked entity represents. You can set a **Flag** color by right-clicking a bookmark row under the **Flag** column, clicking the **Flags** item in the context menu that appears, and then selecting the **Flag** color in the context submenu.
- **Link** icon — this sortable column enables you to attach one or more files to a bookmark, as described in Adding Attachments and Comments.
- **Comment** icon — this sortable column enables you to add one or more comments to a bookmark, as described in Adding Attachments and Comments.
- **Pattern** icon — becomes active when you are adding a bookmark to a **Pattern Group** row. This occurs after you select a pattern match in the **Matched Instances** section of the **Pattern Match** viewer and then select the **Pattern Bookmark** command from the **Add** drop-down menu on the toolbar of the **Bookmarks** window.
- **Scope** — this column is set to **Shared** by default, which enables all users to see the bookmarks you configure. However, in a future Message Analyzer release, you may be able to limit the scope to **Private** so that only you can view the bookmarks you configure in a particular set of trace results.

Using the Bookmarks Toolbar Commands

The commands that are available on the **Bookmarks** toolbar consist of the following:

- **Add** drop-down menu — provides the commands that enable you to configure a **Message Bookmark** or a **Pattern Bookmark**. The enabled/disabled state of these commands corresponds with the session viewer that is in focus.
- **Delete** — enables you to delete a bookmark that you select/highlight in the **Bookmark** configuration grid.
- **View** drop-down list — provides list options that enable you to display bookmarked messages or matched instances in different views:
  - **Analysis Grid** — select this option to display the messages in a **Selected messages** group in a separate instance of the **Analysis Grid** viewer, for analysis of message details and stack information.
  - **Message Gantt** — select this option to display the messages in a **Selected messages** group in a separate instance of the **Gantt** viewer, to correlate the message time range and IP address pair information of the book marked messages.
  - **Pattern Viewer** — select this option to display the **<n> message(s)** locator button in the **MATCHES** pane of the **Pattern Match** viewer for a **Pattern Group** bookmark that you selected in the **Bookmarks** configuration grid. By clicking the message locator, you can display the pattern match instance that corresponds with the **Pattern Group** bookmark that you selected in the **Bookmarks** configuration grid. This is the method you will use to redisplay previously book marked pattern match instances.

- **Pattern Gantt** — select this option to display the messages of a **Pattern Group** in a separate instance of the **Gantt** viewer, to correlate the message time range and IP address pair information of the pattern matched messages.

Adding Bookmarks to an Existing Bookmark or Group

After you create a **User Bookmark** for one or more messages, or create a **Pattern Group** bookmark with one or more messages of a **MATCHED INSTANCE** from the **Pattern Match** viewer in it, you can continue to add bookmarked messages to either of these entities. You can do this by selecting the **Message Bookmark** item in the **Add** drop-down menu to display a submenu that lists any existing **User Bookmark** or **Pattern Group** bookmarks. By selecting an existing bookmark or group in this submenu, you will add any message that is highlighted to the selected bookmark or group. This is true whether you are adding highlighted messages in the **Analysis Grid** viewer or the **Pattern Match** viewer.

You can use this feature to add messages to a bookmark in any conceivable combination. For example, if you are working in the **Pattern Match** viewer, you can add the messages of a selected **MATCHED INSTANCE** to an existing **Pattern Group** bookmark. You could also add one or more messages that are highlighted in the **Analysis Grid** viewer to an existing **Pattern Match** group, or you can add the messages of a **MATCHED INSTANCE** to an existing **User Bookmark**.

Adding Attachments and Comments

To attach a file or add a comment to a bookmark, you must first click the book icon in the **Name** column of the **Bookmarks** configuration grid for the **Selected message** or **Pattern Group** row where you want to attach a file or add a comment. When you click the book icon (or double-click any non-editable field), a drop-down displays with the following three tabs, from where you can perform the indicated tasks:

- **Messages** — the default tab that enables you to view all message rows where you added a bookmark, which includes a display of message numbers and a corresponding **Summary** description for each message in the selected **Bookmarks** message row. Whether you click the book icon for a **Selected messages** row or **Pattern Group** row, the **Messages** tab displays the same type of information.
- **Links** — this tab enables you to add one or more file attachments to a bookmark. You can do this by clicking the open folder icon on the **Links** tab to launch the **Open** dialog and navigate to a file you want to attach; then click the **+** button (green icon with **Add Link** tooltip) to add the attachment to the attachment list. You can also delete any existing attachment by selecting it in the attachment list and clicking the delete (**X**) button.
- **Comments** — this tab enables you to specify one or more comments for any **Bookmarks** message row by clicking the **+Add** button and entering an optional comment **Title**, descriptive text for the comment, and saving the comment by clicking the **Save Changes** icon. You can also delete any existing **Comment** for a selected **Bookmarks** message row by clicking the delete (**X**) button.

> **TIP**
> You can also configure a comment from the **Comments Tool Window**, which contains the identical interface components for configuring comments that you use in the **Bookmarks** window. However, the **Comments** window enables you to create comments that are independent of bookmark configuration.

- **Patterns** — this tab enables you to view all the **Pattern Group** rows where you added bookmarks. This tab contains a header that specifies the following information for each row of data:
  - **Pattern** — the name of the Pattern expression that you executed.
  - **Instance** — the identifying number of the matched instance you selected for a bookmark.

- **Messages** — the number of messages that exist in the selected **Pattern Group** that displays in a **Bookmarks** configuration row.

> **NOTE**
> The book icon in each data row in the **Bookmarks** configuration grid is enhanced with certain glyphs whenever a configured bookmark has an attached file and/or a specified comment. These glyphs consist of the same icons that delineate the **Link** and **Comment** column headers in the **Bookmarks** configuration grid, respectively.

Using Context Menu Commands

The **Bookmarks Tool Window** contains a right-click context menu where you can access several commands that perform the actions indicated below when selected:

- **View Message Range in 'Analysis Grid'** — enables you to display the range of messages that a bookmark contains in a separate instance of the **Analysis Grid** viewer.
- **Add Message Range to Filter** — enables you to automatically create the code for a view **Filter** based on the messages contained in a bookmark row. When applied, the filter will remove all messages from the current **Analysis Grid** viewer instance, except the messages that fall within the range of the selected bookmark. Note that for best results, the **Analysis Grid** viewer that is in the same session as the **Pattern Match** viewer must have the focus.

> **TIP**
> This makes it convenient to save only the messages that were bookmarked by clicking the **Save As** item in the Message Analyzer **File** menu to open the **Save/Export Session** dialog.

- **Delete Bookmark** — enables you to delete any selected bookmark in the **Bookmarks** configuration grid.
- **Show Details** — enables you to expand the details of a selected message row as a drop-down in the **Bookmarks** window. This action also occurs when you click the book icon in a message row.
- **Flags** — enables you to choose a flag color to define a critical issue or level of importance that a particular set of bookmarked messages represents.
- **Open Links** — enables you to open one or more files attached to a **Bookmarks** message row. For example, if an attachment is a trace file, the message contents of the file will display in a separate **Analysis Grid** viewer session tab. If an attachment is an image, it will open in the default image viewer on your system for the particular file type.
- **Copy Selected Rows** — enables you to copy the text in the data columns of a selected **Bookmarks** data row.
- **Copy 'columnName'** — enables you to copy the name of certain bookmark columns, for example, the **Name** and **Category** columns. The single-quoted text in this command is a placeholder for the actual selected column.
- **Edit 'columnName'** — enables you to specify and edit the name of certain bookmark columns, for example, the **Name** and **Category** columns. The single-quoted text in this command is a placeholder for the actual selected column.

Saving and Loading Trace Files Containing Bookmarks

If you want to save bookmarks that you have configured in a message collection, you can only do so by saving

your messages in the Message Analyzer native .matp format. When saving a message collection that contains bookmarks, note that if you use the **Export** option in the **Save/Export Session** dialog to save your message collection to a .cap file, bookmarks will not be saved. When you reload a saved message collection containing bookmarks, for example, through **Open**, **Recent Files**, or a Data Retrieval Session, the bookmarks you configured will display in the **Bookmarks** window, which should automatically open when messages are loaded.

> **IMPORTANT**
> If you are adding bookmarks to a file that you loaded into Message Analyzer through a Data Retrieval Session, you must click the **Save** button on the global Message Analyzer toolbar to save your bookmarks. Otherwise, they will not appear in such a file when you reopen it.

Viewing Bookmarked Data

To view bookmarked messages, load a saved file that contains bookmark data into the **Analysis Grid** viewer. Also, if you expect to be viewing book marked messages associated with **MATCHED INSTANCES** of the **Pattern Match** viewer, you should display the **Pattern Match** viewer from the **New Viewer** drop-down list that is accessible from the global Message Analyzer **Session** menu or on the global toolbar. Also ensure that the **Bookmarks Tool Window** — which is accessible from the **Windows** submenu in the global **Tools** menu — is also displayed.

Thereafter, you can use the following methods to view bookmarked messages:

- **Selected messages drop down** — with the **Analysis Grid** viewer in focus, highlight any **Selected messages** row in the **Bookmarks** window and then click the book icon for the selected row to display a drop-down containing your bookmarked message/s. Select each message in the drop-down to drive selection of the associated bookmarked messages in the **Analysis Grid** viewer.
- **Pattern Group message locator** — with the **Pattern Match** viewer in focus, highlight any **Pattern Group** row in the **Bookmarks** window and then select the **Pattern Viewer** item in the **View** drop-down list on the **Bookmarks** toolbar to display the message locator button in the **MATCHES** pane of the **Pattern Match** viewer. Then do the following:
  - Click the message locator button to display the pattern match instance (associated with the selected **Pattern Group** row) in the **MATCHED INSTANCES** pane of the **Pattern Match** viewer.
  - Select the matched instance that displays and observe the display of messages associated with the bookmark in the **Messages** section of the **Pattern Match** viewer.
  - Select each of these messages to interactively drive selection of each message in the **Analysis Grid** viewer for further analysis, for example, to examine message **Details**.

    > **NOTE**
    > To provide a more convenient view of message selection correlation, undock the **Pattern Match** viewer by dragging it by its tab away from the default location, and then use the docking navigation controls to reposition it directly adjacent to the **Analysis Grid** viewer.

    More Information
    To learn more about the docking features, see Working with Message Analyzer Window Layouts.

- **Pattern Group drop down** — click the book icon to the left of the **Pattern Group** of interest in the **Bookmarks** window to display a drop-down that contains the messages associated with a particular bookmarked **MATCHED INSTANCE** in the **Pattern Match** viewer.

You can then click each message in the drop-down to interactively drive selection of corresponding messages in the **Analysis Grid** viewer. As previously described, you can undock and reposition the **Pattern Match** viewer for a better view of message selection correlation.

With these methods, you or another user can navigate through all bookmarked messages and pattern match instances.

Comments Tool Window

4 minutes to read

Message Analyzer provides the **Comments Tool Window** to enable you to add a comment to any message that you select in any **Analysis Grid** viewer instance. This includes filtered views that display a subset of a full set of trace results.

You typically configure comments to contain important reminders or other information snippets that are significant to data analysis, so that you or others can quickly locate them when resuming data analysis. Therefore, when sharing a trace with commented messages, it will be easier for others to search for and focus on any issues you have flagged with comments. To configure a comment, you can access the **Comments** window as an item in the **Windows** submenu that is accessible from the global Message Analyzer **Tools** menu.

Configuring a Comment

You can add a single comment to any message that you select in the **Analysis Grid** viewer. You can also add multiple comments to a single message that is selected in the **Analysis Grid** viewer, which includes both top-level and child messages. However, if you add one or more comments to a *group* of selected messages, the comment/s will be attached to the last selected message only. After you display the **Comments** window, select a particular message in the **Analysis Grid** viewer that you want to comment. You can then configure the comment by using the following features of the **Comments** window in the indicated manner:

- **Add** — click the **Add** button on the toolbar of the **Comments** window to display the comments input configuration. To configure multiple comments, click the **Add** button successively as required.
- **Title** — enter a title for your comment in this text box.
- **Author** — your user name and the date-time for the comment entry are specified in this text box by default, but you can change it as required.
- **Comment** — specify the text of your comment in the text box below the **Author** text box.
- **Save Changes** — clicking the **Save Changes** icon enables you to save the specified comment configuration.
- **Cancel Edit** — clicking the **Cancel Edit** icon enables you to remove the currently specified comment configuration. If you click the **Cancel Edit** icon, it changes to the **Delete Comment** icon and the **Save Changes** icon changes to the **Edit Comment** icon.

  At this point, you can either resume editing or delete the comment configuration altogether. If you click the **Delete Comment** icon, the **Delete Comment** dialog displays to confirm the deletion. If you click the **Edit Comment** icon, an empty comment configuration displays to enable you to create your comment configuration as required. However, if you happen to have an existing comment selected, clicking the **Edit Comment** button opens the comment configuration for editing.

Searching for Commented Messages

To search for messages that contain configured comments, click the **Previous** and **Next** directional search controls against a set of trace results to find them. You can also use the keyboard left or right arrow keys to select either the **Previous** or **Next** search control and then use the keyboard `ENTER` key successively to search a trace for comments. Note that for large traces, the searches might not be instantaneous. Also, when the search in a specified direction reaches the last message containing a comment, the search simply resumes in round-robin

fashion, should you continue searching in the current direction.

> **NOTE**
>
> If you added a comment to a child message and the origins tree containing the child message is not expanded in the **Analysis Grid** viewer while you are searching for commented messages, only the top-level message or operation containing the commented child message will be highlighted in the **Analysis Grid** viewer. Otherwise, if the child messages are in the expanded state in the **Analysis Grid** viewer, each child message in the stack will be highlighted in succession until the message is reached where the comment was originally configured.

Saving and Loading Trace Files Containing Comments

If you want to save comments that you have configured in a message collection, you can only do so by saving your messages in the Message Analyzer native .matp format. You can save a message collection that contains comments by using the **Save/Export Session** dialog; however, you must click the **Save As** button rather than the **Export** button to facilitate the save; otherwise, your comments will not be saved in a .cap file which is the result of saving with the **Export** command. When you reload a saved message collection containing comments, for example through **Open**, **Recent Files**, or a Data Retrieval Session, the **Comments** window automatically opens when Message Analyzer is finished loading the data. Thereafter, you can use the **Previous** and **Next** buttons in the **Comments** window to locate the associated commented messages in the **Analysis Grid** viewer. In this way, you or another user can navigate through all of your commented messages.

Other Windows

2 minutes to read

This section describes several other **Tool Windows** that you will make regular use of in Message Analyzer, which include the following:

- **Field Chooser Tool Window** — provides message hierarchies in tree grid format to assist in configuration tasks where field selection is required, for example, when creating **Pattern** expressions, **Charts**, and **Unions**, or when adding new data columns to the **Analysis Grid** viewer default column layout or new Groups to the **Grouping** viewer.
- **Output Tool Window** — provides the output of the Message Analyzer log, as a convenience, to enable you to monitor Message Analyzer start up routines such as loading parser modules. Any errors that occur are highlighted in red in this window, which is useful information if you want to file a bug (click **Report a bug** in the **Feedback** drop-down menu). Also displays other data such as .cap file export statistics.
- **Session Explorer Tool Window** — enables you to navigate among session and viewer nodes in Live Trace Session and Data Retrieval Session results, provides a context menu for opening new data viewers, and also displays session progress indications and other statistics.
- **Map Tool Window** — provides a high-level view of data points along with a navigation control that enables you to focus on the data field locations that contain the messages you want to analyze in a supporting primary data viewer that contains map style visualizations.

You can display any of these windows by selecting them in the **Windows** submenu that is accessible from the global Message Analyzer **Tools** menu, if they are not already displayed. Note that the **Field Chooser** window is also accessible from the **Pattern Editor**, the **Edit Union** dialog, and when creating and editing **Charts**.

The indicated **Tool Windows** are described in the following topics of this section:

Field Chooser Tool Window
Output Tool Window
Session Explorer Tool Window
Map Tool Window

Field Chooser Tool Window

12 minutes to read

The **Field Chooser Tool Window** is a single-instance and dockable window that enables you to specify message types, field types, methods, properties, annotations, and other general fields to enhance the functionality of various Message Analyzer data viewers, for example, as described in Expanding the Analysis Grid Viewer Column Layout. The **Field Chooser** is a common dialog that presents expandable hierarchical nodes that contain the properties, fields, methods, and other entities, for the protocol/module messages that Message Analyzer parses. Note that **Field Chooser** is not an interactive window since it does not drive message or data selection, nor is it driven by such selections.

Displaying the Field Chooser
You can access the **Field Chooser** window from any of the following locations in Message Analyzer to open and dock the window in its default location, if it is not already displayed. If the **Field Chooser** window is already displayed, the following actions will simply place focus on it:

- **Analysis Grid** group — when an **Analysis Grid** viewer session tab has focus, click the **Add Columns** button on the **Analysis Grid** viewer toolbar.
- **Grouping** group — when the **Grouping** viewer has focus, click the **Add Groupings** button on the **Grouping** viewer toolbar.
- **Windows** submenu — click the global Message Analyzer **Tools** menu, click the **Windows** item, and then select the **Field Chooser** item from the **Windows** submenu.
- **Column context menu** — right-click an **Analysis Grid** viewer column header and select the **Add Columns...** item that displays in the context menu that appears.

Redocking the Field Chooser
After you display the **Field Chooser** window, you can undock and reposition it by taking advantage of the docking navigation control that displays after you drag the **Field Chooser** window away from its default docking location. You might do this to configure a location for this window that better suits your analysis process.

Performing Tasks with the Field Chooser
You can use **Field Chooser** to locate message fields when performing the following tasks:

- Expanding the default column **Layout** of the **Analysis Grid** viewer by adding fields as new data columns, as described in Expanding the Analysis Grid Viewer Column Layout.
- Adding new groups to the **Grouping** viewer by specifying the fields that you want to become the new Groups, as described in Expanding the Grouping Viewer Configuration.
- Adding configuration settings to various data viewers and tools by specifying fields for **Chart** viewer **Layouts**, **Pattern** expressions, **Unions**, and so on, as described in Locating Message Fields for Configuring Settings.

Expanding the Analysis Grid Viewer Column Layout

The **Field Chooser** window enables you to configure new data columns for the **Analysis Grid** viewer to display values for the fields of specified protocols and modules. The default column **Layout** of the **Analysis Grid** viewer provides a starting point for data analysis; however, the default **Layout** offers a limited cross section of more robust data sets that are available to you through use of the **Field Chooser** window. For

example, your data analysis process might be focused on a particular protocol for which the **Analysis Grid** viewer default column **Layout** displays only a subset of the greater superset of field information that you need to examine. With the **Field Chooser** window, you can add the columns you need to display the values of any data field you want to view, which includes message types, properties, structures, methods, flags, events, or other fields that a protocol of interest defines.

When you open the **Field Chooser** window, it displays a tree view of nodes that represent protocols and modules for which Message Analyzer provides parsing based on OPN descriptions, in addition to various annotations, properties, global fields, and any **Unions** that you have created. To access the fields of a particular protocol or module in the **Field Chooser** window, you can expand its top-level node and subnodes as required to expose the message hierarchy containing the message types, properties, structures, methods, flags, and events that the protocol defines.

Adding and Removing Analysis Grid Data Columns

The specific data columns that you will add to the **Analysis Grid** viewer during data analysis generally depends on the troubleshooting context in which you are engaged. For example, prior to running a trace, you might have clients who complain about an unresponsive or slowly responding web server. After you run a trace that targets the server traffic and Message Analyzer displays the data, you might notice that the default column layout of the **Analysis Grid** viewer provides some initial indications in the **Summary** column about HTTP requests and responses. However, you might want to add HTTP **StatusCode**, **ReasonPhrase**, **HTTPContentType**, and **Method** columns, in addition to the **ResponseTime** Global Annotation as a column, so you can **Group** these columns to better organize and expose the information for analytical purposes.

To add a new data column to the **Analysis Grid** viewer and expose the data values for a field of any protocol or module that Message Analyzer parses, you will need to navigate through the appropriate message hierarchy in the **Field Chooser** window until you find the required field name/s. You can also search for a field by name by entering it in the search textbox, at which time search results are highlighted for any matches that are found. At this point, you can simply select the field you want to add as a column to the **Analysis Grid** viewer and then click the **Add** button on the **Field Chooser** tool bar. You can also right-click the field you want to add and select the **Add as Column** context menu item. Note that you will be unable to use the **Add as Column** command unless an **Analysis Grid** viewer tab has focus; otherwise, this command will be disabled.

> **TIP**
>
> You can remove any data column from the **Analysis Grid** by right-clicking the column you want to delete and then selecting the **Remove** item from the context menu that appears. From this context menu, you can also select the following commands to perform related operations after adding new fields to the **Analysis Grid** viewer with **Field Chooser**:
>
> - **Save as Default User View Layout** — enables you to save any column **Layout** configuration as the user default **Layout**.
> - **Load Default User View Layout** — enables you to load the column **Layout** configuration that you saved as the user default **Layout**.
> - **Save Current View Layout As...** — enables you to save the current column **Layout** configuration with a unique **Name**, **Description**, and **Category** specification from the **Edit Item** dialog.

Working with Array Fields

You might notice that some fields in **Field Chooser** are denoted as arrays **[]**, which require you to specify an array element or **Key** value in order to see the data in a new column that you are adding to the **Analysis Grid** or **Grouping** viewer. For example, if you wanted to look at the data for a **Uri.Query**, first expand the **Uri** node in the **HTTP** message hierarchy in **Field Chooser** then expand the **Query** node. Immediately under the **Query** node, you will see an array [] designator. If you right-click this designator and select **Add as Column**, the **Collection Key Selector** dialog displays to enable you to specify a **Key** such as "ocid" or some other string (you can correlate these in the **Details Tool Window** for the selected message). After you click **OK** to exit the

dialog, the data for the **Key** you specified will display in the **Analysis Grid** under a new column entitled **Uri.Query["ocid"]** in this example.

A simpler method for doing this would be to locate the array value for which you want to display data in the **Details Tool Window**, right-click it, and select the **Add as Column** command. However, note that the formerly specified method is most useful when **Details** for a particular field are unavailable for a particular message. An example might be that you want to look at SACK options for TCP messages, but it is difficult to find any TCP messages that expose them such that you could right-click a field in **Details** and add a new column in the **Analysis Grid** viewer. When this is the case, you would need to be very familiar with TCP **Options** and values so that you could specify the required **Key** value/s in order to view the data.

More Information
To learn more about these commands in addition to saving, sharing, and updating view **Layouts**, see Applying and Managing Analysis Grid Viewer Layouts.

Expanding the Grouping Viewer Configuration

While the **Grouping** viewer has focus, you can expand the Group configuration by adding one or more fields to the current **Grouping** viewer **Layout**. To do this, click the **Add Groupings** button on the **Grouping** viewer toolbar to cause the **Field Chooser** window to open in focus or to set the focus if it is already open, enabling you to locate fields of choice to add to the nested Group configuration of the **Grouping** viewer for an enhanced analytical perspective. For example, after you navigate to a protocol node of interest in **Field Chooser**, you can add a new Group to the **Grouping** viewer by selecting the **Add as Grouping** context menu item that displays when you right-click a particular field that you select under the current message hierarchy. With this capability, you can extract additional information into the Group configuration to expose data that might be difficult to locate in a large data set, while also enhancing the interactive analysis context with other data viewers, as described in the Grouping Viewer topic.

Locating Message Fields for Configuring Settings

You can use the **Field Chooser** when you are creating new **Pattern** expressions, new **Chart** viewer **Layouts**, new **Grouping** viewer **Layouts**, and **Unions**. The **Field Chooser** window enables you to specify the fields of protocols and modules as configuration settings when modifying viewer or other entity configurations. To do so, you can navigate through message hierarchies to find fields of interest in the same way that you do when adding columns to the **Analysis Grid** viewer.

In this context, you can display the **Field Chooser** window when using the following features in the indicated ways:

- **Pattern** expression configuration — while an **Analysis Grid** viewer tab has the focus, right-click the corresponding session in the **Session Explorer Tool Window**, select **New Viewer**, and then select the **Pattern Match** item from the context menu that displays. After the **Pattern Match** viewer displays, open the **Pattern Editor** by clicking the **Create Pattern** button on the **Pattern Match** viewer toolbar. You can then display the **Field Chooser** window by clicking the **Insert Message** button in the lower section of the **Quick** tab in the **Pattern Editor**.

  At this point, the **Field Chooser** window enables you to specify a message type for the **Pattern** expression configuration. Thereafter, the **Pattern Editor** enables you to again display the **Field Chooser**, which opens to the fields, properties, methods, and so on that are associated with the previously specified message type. You can do this by first clicking the **Insert Criteria** link on the **Quick** tab and then by clicking the ellipsis (**...**) control under the message configuration toolbar.

- **Chart** viewer **Layout** configuration — after you have the **Chart** viewer with a particular **Layout** displayed, you can then create a new **Chart** or customize the current one. To begin, click the **Edit** item in

the **Chart** menu that displays when you select **Chart** in the global Message Analyzer **Session** menu. The **Edit Chart Layout** dialog then displays from where you can specify a **Chart** visualizer component that you want to use, for example, a **Bar** chart or a **Table** grid that you select in the **Chart type** drop-down. Thereafter, you can specify data fields and create formulas for the rows and/or columns of the **Chart** visualizer component.

You can display the **Field Chooser** to specify the data fields by clicking the ellipsis (...) in the **Series** pane of the **Edit Chart Layout** dialog. You will also need to display **Field Chooser** when configuring a data formula by clicking the ellipsis in the **Values** pane of the **Edit Chart** dialog, and once again in the **Formula Editor** that displays immediately thereafter. See Extending Message Analyzer Data Viewing Capabilities for further details.

- **Grouping** viewer **Layout** configuration — while the **Grouping** viewer has focus, click the **Add Groupings** button on the **Grouping** viewer toolbar to display the **Field Chooser** window, as described earlier in Expanding the Grouping Viewer Configuration.
- **Union** configuration — after you click the **New Union** button on the global Message Analyzer toolbar, the **Edit Union** dialog displays. When you click the **Add** button in this dialog for fields to include in the **Union**, the **Field Chooser** window displays to enable you to locate and select the fields you want in the **Union** configuration.

Understanding Other Field Chooser Categories

The **Field Chooser** window contains a number of nodes that are not directly associated with a protocol or module; however, you can use the entities contained in these nodes in any of the tasks described in this section, among many others. These nodes are located in the uppermost categories of **Field Chooser** window and consist of the following:

- **General** category — contains common entities that you can add as new columns in the **Analysis Grid** viewer, use in a **Filter**, configure in a **Pattern Expression**, and so on. You can also use **Global Annotations** and **Global Properties** in these tasks and many others. Some entities in the **General** category consist of **Type**, **DiagnosisLevels**, and **DataSource**.
- **Global Annotations** category — additional information that is not directly related to a message, such as user-provided comments, implementation data, or other information related to the network stack. To retrieve this information, for example in a Filter Expression, you must use the number symbol operator ("#"). For example, if you wanted to locate a message displayed in the **Analysis Grid** viewer, you could use a Filter that is similar to the following:
  ```
  #MessageNumber == 5
  ```
- **Global Properties** category — contains a group of properties that Message Analyzer provides to enhance functionality. Some of the properties are inherent to existing message types, while others were created specifically for Message Analyzer to expand the data that is available to you for analysis. For example, the property **IsOperation** can identify whether a message node in the **Analysis Grid** viewer is an Operation. This can enable you to return Operation messages only with the use of a Filter Expression such as the following:
  ```
  *IsOperation == true
  ```

  Please be aware that you must use an asterisk (*) to return the data of a **Global Property**.

Note that the **Global Annotations** and **Global Properties** entities appear in the **Details Tool Window** when you click the **Show all properties for the selected message** button on the **Details** window toolbar. However, these entities will display a value in **Details**, only if such values can be read or otherwise derived from the currently selected message.

More Information

To learn more about using the **Field Chooser** with various Message Analyzer features, see the following topics:

Using the Field Chooser

Accessing Message Properties and Annotations

Using the Pattern Editor

Extending Message Analyzer Data Viewing Capabilities

Grouping Viewer

Creating Unions

Output Tool Window

2 minutes to read

The **Output Tool Window** is a dockable window that outputs text from the Message Analyzer log file as a result of various operations, for example, when Message Analyzer loads modules at startup. If errors are logged, the **Output** window will show them in red text. As a usage example, the **Output** window can conveniently expose when an OPN parser fails to load. Another example is that the **Output** window displays statistical data when you export a trace to the .cap file format from the **Save/Export Session** dialog, which is accessible by clicking either the **Save** or **Save As** item in the Message Analyzer **File** menu.

> **NOTE**
>
> If the **Output** window logs errors, you can use this information to file a bug by clicking the **Report a bug** item in the **Feedback** drop-down menu.

You can display the **Output** window by selecting it from the **Windows** submenu that is accessible from the global Message Analyzer **Tools** menu.

Session Explorer Tool Window

5 minutes to read

The **Session Explorer Tool Window** is a single-instance, interactive, and dockable window that drives the selection of Message Analyzer session viewers that can contain data from Live Trace Sessions and/or Data Retrieval Sessions in different view configurations. The **Session Explorer** window also provides a **New Viewer** context menu that enables you to select numerous types of data viewers and **Layouts** to assist your data analysis process.

The **Session Explorer** window configuration contains session nodes and viewer subnodes, along with associated progress indicators that display when messages are being loaded, captured, or processed by Message Analyzer. All viewer subnodes in the **Session Explorer** window are contained under top-level session nodes, and can consist of one or more viewers for each Live Trace Session and/or Data Retrieval Session in display.

Whenever you have session data for which there are multiple data viewers displayed, you can bring the data contained in a particular viewer into focus by clicking the appropriate viewer subnode in the **Session Explorer** window. For example, a session node for a Live Trace Session might have **Analysis Grid** and **Gantt** session viewer subnodes under it that represent the corresponding data viewers that you opened for that session. If you successively select these subnodes, the selection action drives the display of data in the **Analysis Grid** viewer tab and then in the **Gantt** viewer tab. Note that you can also bring the data of a particular viewer into focus by simply selecting its session viewer tab in an Analysis Session.

Using Session Explorer Features

The functions and features of the **Session Explorer** window enable you to do the following:

- **Select new session data viewers** — **Session Explorer** has a context menu that displays when you right-click any session node, viewer subnode, or anywhere in the white space of the **Session Explorer** window. The **New Viewers** command in the context menu opens a submenu that provides a selection of numerous data viewers that you can choose to enhance your data analysis perspectives. For some viewers, such as the **Analysis Grid**, **Grouping**, and **Chart** viewers, you can specify a **Layout** from a drop-down list that displays when you highlight one of these viewers in the indicated submenu.

  From the **New Viewers** submenu, you can select any of the data viewers that are described in the Data Viewers section of this documentation. After you select one or more viewers, Message Analyzer displays data from the current session in a separate viewer tab for each different type of viewer.

  The **Session Explorer** context menu also has a **Close** command that closes the session represented by a selected top-level session node, along with all open viewers in such a session. Note that if you have more than one data viewer open in a particular session, you can close a single data viewer only, without closing the session, by right-clicking its node and selecting the **Close** command.

- **Navigate session viewer tabs** — after you open different data viewers for one or more sessions and data displays in separate session tabs, the **Session Explorer** window enables you to quickly navigate through the data contained in each session viewer tab. For example, by selecting any viewer subnode in **Session Explorer**, you can instantly display the session viewer tab containing the data that corresponds to the selected viewer subnode. By polling through your data viewers and presentation formats in this manner, you achieve different data perspectives that can enhance your analysis experience.

  In addition, for each session that you open and display data in a selected Message Analyzer viewer, a color-coded dot is configured in the top-level session node in **Session Explorer**, for ease of identification. Thereafter, all data viewers that you open in an identical session will also contain the same color-coded dot

in its viewer node in **Session Explorer**, thus automatically organizing the applicability of your data viewers. Note that the same color code is also applied to the associated session viewer tabs that display in the data analysis surface of the Message Analyzer UI.

- **Observe session and viewer tooltip information** — if your mouse hovers over a top-level session node in **Session Explorer**, a tooltip displays to indicate the number of messages in the session. Also, if your mouse hovers over a session viewer node in **Session Explorer**, you can see the analytical assets that are applied to the associated session viewer. For example, the tooltip can display the currently applied filtering configuration, which can include a **Viewpoint**, **Viewpoint Filter**, **Time Filter**, **Message Range Filter** (from the **Gannt** viewer), and/or view **Filter**.

- **Monitor session performance and status** — you can monitor session performance by observing *relative* progress indicators that display when you are loading saved trace and log file data into Message Analyzer through a Data Retrieval Session, or you can observe *ambiguous* progress indicators when you are capturing data in a Live Trace Session. Relative progress indications are also provided whenever you apply data manipulation functions to a set of messages, which includes filtering, sorting, finding, grouping, and pattern matching. Relative progress indicators are more definitive because Message Analyzer can determine the number of messages being loaded; whereas for ambiguous progress indicators, this is not the case.

  You can also monitor session statistics such as the following in the Message Analyzer status bar area:

  - **Session status** — for example, you might see an indication such as "Ready" or "Processing".
  - **Session Total** — the total number of messages in a session.
  - **Available** — the number of available messages in a session after operations such as **Viewpoints** and view **Filters** are applied to a set of trace results.
  - **Selected** — the number of messages that are currently selected in a set of trace results.
  - **Viewpoint** — the current **Viewpoint** that is applied to a set of trace results.
  - **Truncated Session** — indicates with a **True** or **False** label whether or not a session has truncated messages.
  - **Parsing Level** — the current **Parsing Level** that is applied to a set of trace results.

Displaying Session Explorer

If the **Session Explorer** window is not displayed, click the global Message Analyzer **Tools** menu, click the **Windows** item, and then select the **Session Explorer** item to restore it. Note that you can also undock and reposition the **Session Explorer** window by taking advantage of the docking navigation control that displays after you drag the **Session Explorer** window away from its default docking location. You might do this to configure a location for the window that is more convenient for the analysis environment in which you are working.

More Information
To learn more about selecting data viewers, see Session Data Viewer Options.
To learn more about progress monitoring and session statistics, see Viewing Session Statistics and Progress.
To learn more about docking and redocking tool windows, see Working with Message Analyzer Window Layouts.

Map Tool Window

2 minutes to read

Message Analyzer provides the **Map Tool Window** to support data analysis in primary viewers that make use of data maps of various types, such as the **Gantt** and **Interaction** viewers. The **Map** window provides a high-level view of data points along with a navigation control that enables you to focus the corresponding primary data viewer on the data field locations that contain the messages you want to analyze. In the **Gantt** viewer, this means that you can quickly locate specific color-coded message groups on which you want to focus during analysis. For the **Interaction** viewer, this means that you can quickly locate data points in the **Mapping**, **Diagram**, and **Chord** display formats of the **Interaction** viewer that contain the messages on which you want to focus your analysis. The modes of operation and zooming capabilities of the **Map** window are described in the section that follows.

Map Tool Window Controls

The **Map** window contains several commands that you can utilize from a context menu that displays when you right-click anywhere in the **Map** window data field. These commands consist of the following:

- **Mouse Mode** — provides the following commands:
  - **Zoom** — enables the **Zoom** mode, where you can create a data selection window by clicking and dragging your mouse across the **Map** window data field.

    The **Zoom** mode works together with options of the **Zoom** item described below, by altering the zooming ratios of the **X** axis, **Y** axis, or both **X and Y** axis together, to enable you to focus on specific messages that you want to analyze. The presets that you can select to change the zoom ratio are accessible from the secondary context menu of the **Zoom** item described below.

> **TIP**
> After you create a data selection window, you can click anywhere in the **Map** window data field and cause the selection window to track your mouse clicks by jumping to that location.

- **Pan** — enables the **Pan** mode, where you can drag a fixed-size data selection window across the **Map** window data field to locate the messages on which you want to focus your analysis.
- **Zoom** — enables you to change the zoom ratio in the **Map** window by selecting preset zooming values for the **X** axis, **Y** axis, or **X and Y** axis together, as follows:
  - **1:1**
  - **Fit**
  - **In**
  - **Way In**
  - **Out**
  - **Way Out**
- **Fit Map to Window** — by selecting this mode, it enables you to fit the entire data field of the in-focus viewer into the **Map** window. By unselecting this mode, the **Map** window data field size no longer encapsulates the data field of the in-focus viewer.

Working with Message Analyzer Window Layouts

16 minutes to read

By default, Message Analyzer provides several built-in **Window Layouts** that organize the **Analysis Grid** viewer with different **Tool Windows** as preset configurations that enable you to customize your working environment for the type of troubleshooting and analysis in which you are engaged. The window layouts that you can choose range from simple to increasingly more complex selections in the **Window Layout** drop-down list, and are intended to accommodate a cross-section of typical Message Analyzer users. You can access this drop-down list from the global Message Analyzer toolbar. The typical layout configuration consists of a single/default data viewer and an arrangement of one or more **Tool Windows**. By default, Message Analyzer uses the **Analysis Grid** viewer in all the built-in **Window Layouts**; however, you can set a different **Default Viewer** from the **Profiles** tab of the **Options** dialog that is accessible from the global Message Analyzer **Tools** menu. You can also add other **Tool Windows** to any of the built-in **Windows Layouts**, as needed.

You also have the option to further customize your analysis environment by redocking any data viewer or **Tool Window** in use to a location that facilitates easier viewing and analysis, or better correlation of related message data that is held in different views, as described in Using the Redocking Features. If you do not select one of the built-in **Window Layouts**, Message Analyzer still provides a default window layout that contains the **Analysis Grid** viewer and several **Tool Windows** that provide a basic configuration for analysis, as described in Using the Message Analyzer Default Window Layout. You will see this default layout configuration at first Message Analyzer startup and upon all subsequent startups unless you change the layout by any of the following actions:

- Manually open other **Tool Windows** from the **Windows** submenu in the global Message Analyzer **Tools** menu.
- Remove one or more **Tool Windows** from the default layout.
- Select one of the built-in **Window Layouts** on the global Message Analyzer toolbar.
- Redock viewers and/or **Tool Windows** to other locations to create a custom configuration for your environment.
- Manually change the **Default Viewer** on the **Profiles** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu.
- Load data from a file type for which a viewer configuration is defined by a default or custom **Profile**, as described in topic Working With Message Analyzer Profiles.

The following topics in this section describe how to work with **Window Layouts**.

Using the Message Analyzer Default Window Layout
Using the Built-In Window Layouts
Redocking Data Viewers and Tool Windows
Performing Docking Operations
Using the Redocking Features
Saving Window Layouts

Using the Message Analyzer Default Window Layout

The default viewer and **Tool Window** layout configuration that displays on the first Message Analyzer startup is briefly described in the list below. You can continue to use the default layout as is, or you can add other **Tool Windows** from the **Windows** submenu of the global Message Analyzer **Tools** menu as needed. The **Tool**

Window configuration that you specify is recorded and persisted through subsequent Message Analyzer startups, until you change it to a new configuration, which in turn is persisted until you change it again. You can also redock the **Analysis Grid** viewer or any **Tool Window** to another location by making use of the Docking Navigation Control that displays when you undock a window and hover over certain user interface (UI) locations. However, the **Analysis Grid** viewer, or whatever viewer is set as the default on the **Profiles** tab of the **Options** dialog, will persist in its default location on Message Analyzer restarts. The default viewer and **Tool Window** configuration that is provided by Message Analyzer consists of the following:

- **Analysis Grid** — a redockable and interactive tree grid configuration display with expandable message nodes that is the default Message Analyzer viewer for data analysis. It includes inline stack, details, and diagnosis information that is accessible from each top-level message; encapsulation of network stack messages; related Operation messages condensed into a common node; along with sorting, grouping, filtering, and so on. Includes a default column **Layout** for common analysis of trace results.

  See the Analysis Grid Viewer topic for further details on analyzing data with this viewer.

- **Message Stack 1** — a redockable and interactive **Tool Window** that enables you to view the network stack layers (origins tree) for any selected message that contains a message stack configuration in a Data Retrieval Session or Live Trace Session.

  See the Message Stack Tool Window topic for further details on how to use this feature.

- **Details 1** — a redockable **Tool Window** containing message field names, values, types, bit offset, and bit length specifications, that is interactively driven by message selection in the **Analysis Grid** viewer and other viewers.

  See the Message Details Tool Window topic for further details on how to use this feature.

- **Message Data 1** — a redockable **Tool Window** that contains hexadecimal values by default for fields that you select in the **Details Tool Window**.

  See the Message Data Tool Window topic for further details on how to use this feature.

- **Field Data** — a redockable **Tool Window** that displays a string or other value for fields that you select in the **Details Tool Window**.

  See the Field Data Tool Window topic for further details on how to use this feature.

- **Session Explorer** — a redockable and interactive **Tool Window** that enables you to open new data viewers or explore and select open sessions and the viewers in each session for data captured or loaded into Message Analyzer.

  See the Session Explorer Tool Window topic for further details on how to use this feature.

- **Field Chooser** — a redockable **Tool Window** that enables you to view all of the fields that Message Analyzer can parse. From the **Field Chooser**, you can use the **Add as Column** and **Add as Grouping** context menu commands to add new columns of data or groups to either the **Analysis Grid** viewer or **Grouping** viewer, depending on which viewer has focus when the chosen command is executed. You can also display the OPN definition for a field by executing the **Go To Definition** context menu command from the **Field Chooser**.

  See the Field Chooser Tool Window topic for further details on how to use this feature.

Other redockable **Tool Windows** that you can manually select for display in any Message Analyzer session are available as items in the **Windows** submenu which is accessible from the global Message Analyzer **Tools** menu. This includes the **Bookmarks**, **Comments**, **Diagnostics**, **Decryption**, **Selection**, and **Output Tool Windows**.

Using the Built-in Window Layouts

The built-in **Window Layouts** that are provided by Message Analyzer are described in the list that follows. You can select these layouts from **Window Layout** drop-down list on the global Message Analyzer toolbar.

- **Simple** — this layout is the simplest configuration that is available from the **Window Layout** drop-down list. It contains the **Analysis Grid** viewer and the **Details Tool Window** for a basic analysis of messages and message field data. After you select a message in the **Analysis Grid**, the message field names, values, types, and so on, display in the **Details** window. The context menus are still available for both the **Analysis Grid** viewer and **Details** window, so you can continue to use the right-click method to create Filters, Groupings, and so on.

  You might use this layout to perform simple analysis of messages and field data, when analysis of the network layers is of less importance.

- **Simple with Field Data** — this layout is also a basic configuration similar to the **Simple** layout, only it adds the **Field Data Tool Window** to the layout. The **Field Data** window displays field values when you select field names in the **Details** window.

  Similar to the **Simple** layout, you might use this layout to perform simple analysis of messages and field data, although this layout also enables you to quickly assess the alphanumeric values of fields that you select in the **Details** window.

- **Network** — this layout adds the **Message Stack Tool Window** to the **Simple with Field Data** layout configuration to facilitate quick analysis of the network message stack, while retaining all the capabilities of the previously described layouts.

  You might use this layout when troubleshooting the Network stack protocols that support top-level operations or other transactions of a particular application.

- **Multiple Sessions** — this layout adds the **Session Explorer Tool Window** to the **Network** layout configuration to enable you to quickly select the data that is presented in any viewer in one or more sessions. This layout also adds the hexadecimal **Message Data Tool Window** to facilitate an additional view of the hexadecimal values of message fields as they exist in the context of an entire message stream. It also provides the **Output** window, which displays text from the Message Analyzer log file that indicates errors and other statistical information.

  You might use this layout when you have multiple sessions and one or more data viewers in each session, to instantly display the data in any data viewer with a single click. This provides the capability to quickly correlate and assess your data in multiple presentation formats for an enhanced analytical perspective.

- **Protocol Development** — this layout adds the **Compare Fields Tool Window** to the **Multiple Sessions** layout, which can be advantageous when using Message Analyzer to validate protocol field values, states, and behaviors in protocol development and testing scenarios. In addition, the **Message Data Tool Window** can be useful for analyzing hexadecimal data and determining whether some message fields are not being parsed. This layout also removes the **Session Explorer Tool Window**, which is typically not required in these scenarios.

  You might use this layout when troubleshooting the client- or server-side communications of a new protocol in development.

- **Advanced** — this layout adds the **Selection**, **Bookmarks**, **Comments**, and **Session Explorer Tool Windows** to the previous **Protocol Development** layout and relocates some **Tool Windows** such as the **Field Data** and **Output** windows.

  You might use this layout when it is important to annotate messages for follow-up data reviews by others, when you need to see Message Analyzer errors and other statistics provided in the **Output** window, and/or to take advantage of the Message Analyzer advanced message **Selection** and tracking capabilities.

NOTE

You have the option to add other **Tool Windows** to any built-in **Window Layout**, as required, and they will be saved and persisted across subsequent Message Analyzer restarts until you modify the configuration again.

Redocking Data Viewers and Tool Windows

Message Analyzer provides you with the flexibility to reposition the different data viewers and **Tool Windows** that contain your trace results information, by undocking them from their default locations and redocking them in another sector of the UI. After you undock a data viewer or **Tool Window**, you have the option to float, resize, redock, or remove it, as described in Performing Docking Operations. You can also restore any **Tool Window** that you previously removed (closed). Message Analyzer provides the capability to reposition any data viewer or **Tool Window** mainly for the enhancement of your analysis perspective, as described in Using the Redocking Features.

Utilizing the Docking Navigation Controls

To facilitate window relocation, Message Analyzer provides the Docking Navigation Control, which is a simple mechanism that enables you to redock any data viewer or **Tool Window** in a new location after you have initially undocked it. This control is built into the UI as part of the window docking infrastructure and displays immediately after you drag any data viewer tab or **Tool Window** tab away from its current docking location. The Docking Navigation Control displays in various UI locations to facilitate the window relocation capabilities; however, the size of the control varies depending on the section of the UI where you drag a window. The Docking Navigation Control contains a central tab and four directional arrow tabs that enable you to preview a drop location whenever you hover over one of these elements while dragging an undocked window with your mouse. As previously indicated, you can use the Docking Navigation Control to configure a new location for a data viewer or **Tool Window** that is more convenient or better suited to your analysis environment.

Performing Docking Operations

There are several operations that you can perform when working with data viewer and **Tool Window** locations. To move a data viewer or **Tool Window** from its default or current docking location, click a data viewer or **Tool Window** tab and drag it away from the docking location. The operations that you can perform when relocating **Tool Windows** and data viewers consist of the following:

- **Float** — you can float a data viewer or **Tool Window** by dragging it from its default docking location, as previously indicated. After you float a data viewer or **Tool Window**, you can redock it, or leave it in the undocked state, either within or outside the Message Analyzer UI, and it will continue to interact appropriately with other data viewers or **Tool Windows**, just as it did from its previous docking location. For example, you can continue to drive the population of data into a floating **Message Data Tool Window** through message selection in a floating **Analysis Grid** viewer.

  If you choose to redock a floating data viewer or **Tool Window** to its default or other preferred location, you can do so by dropping the window on an appropriate directional arrow of the Docking Navigation Control.

- **Resize** — you can resize any **Tool Window** or data viewer, floating or otherwise, in the same way that you resize any window, which is by hovering over any window edge with your mouse and using the expansion/contraction arrow to select and guide the direction of the resize.

- **Redock** — you can redock any window, floating or otherwise, by dragging it with the mouse and positioning it over one of the directional arrows of the Docking Navigation Control that previews the location you want. When you have selected the new location, release the mouse button.

- **Remove** — to remove (close) a data viewer or **Tool Window** in its current display location in a Message Analyzer session, click the **X** mark on the viewer session tab or on the title bar of the **Tool Window**. If a

data viewer or **Tool Window** is floating, you can remove it the same way.

Using the Redocking Features

When using Message Analyzer to analyze data, a typical scenario might be to display the results of several different sessions in different data viewer formats. It is also common to display the results of a single session in multiple data viewer formats. However, Message Analyzer data viewers such as the **Analysis Grid** are by default positioned in a row of session tabs, one for each viewer instance, where only the data of the in-focus tab is visible. To view the data in another session viewer tab, you must select that tab. But at times, it may be advantageous to expose the data of multiple viewers simultaneously, to achieve an interactive context for comparative purposes when you are correlating data.

Driving Interactive Display of Data

Taking the previous scenario as an example, you might want to interactively drive message displays in the **Analysis Grid** viewer by choosing certain **Matched Instance** messages from the **Pattern Match** viewer that met the filtering criteria of a particular **Pattern** expression and display them in the **Analysis Grid** viewer. In this configuration, or by interactively driving the **Analysis Grid** viewer display from message selection in another viewer instance such as the **Gannt** viewer, you can immediately assess the context of the points in time where the associated messages occurred in a trace.

To position the viewers in this configuration so you can simultaneously and interactively view the data in these different formats, you can undock the **Pattern Match** viewer (or **Gannt** viewer) and redock it to display next to the **Analysis Grid** viewer, so that the data in both viewers is visible at the same time, but in separate session tabs. Note that you have the option to float the viewer you are undocking while you assess where your new docking location will be. While dragging the window you are relocating over the Docking Navigation Control tab or arrows, release your mouse button when the control previews the docking location you want.

After the redock is complete, you can select any **Matched Instance** in the **Pattern Match** viewer and the messages of the **Matched Instance** immediately display in the **Analysis Grid** viewer. A similar type of interaction will occur if you select a color-coded message bar in the **Gannt** viewer. The advantage of seeing these messages in the **Analysis Grid** viewer, is that you can quickly examine the details, field values, and stack configuration of messages in this view.

Another useful docking configuration that Message Analyzer provides is the default location of the **Grouping** viewer and the **Analysis Grid** viewer. In the default locations of these viewers, you can easily correlate data in any filtered group that you select in the **Grouping** viewer with corresponding message data that displays in the **Analysis Grid** viewer, that is, when the **Grouping** viewer is in the **Selection** mode. However, you can still move the **Grouping** viewer to a new docking location in accordance with your preferences.

Correlating Trace and Log Data

At other times, you might need to compare different traces and related log file data, perhaps from different time zones, or you might want to review historical versus current trace results. When this is the case, you can follow the processes outlined earlier in this section to relocate your data viewers in appropriate positions.

> **TIP**
>
> When you are comparing trace results, you can take advantage of the **Compare Fields Tool Window**. This feature enables a quick comparison of the field values of any two messages, as described in the Compare Fields Tool Window topic.

> **NOTE**
>
> If you have many data viewers displaying in multiple sessions, the session tabs that contain viewer data can exceed the width of the available space in the Message Analyzer UI. When this is the case, you can scroll to the hidden session tabs by clicking the direction arrows (**< >**) in the upper-right corner of the session tab row. In addition, if the position of a viewer appears to be locked in place when you try to drag an edge to reposition it, you can use the direction arrows to move the viewer window position while the particular viewer is in focus.

Saving Window Layouts

Message Analyzer automatically saves the current window layout when you shut the application down, which can be one of the built-in **Window Layouts** that are accessible from the global Message Analyzer toolbar or a custom window layout that you configured for your environment. In either case, the window layout that existed at the time of shutdown is then loaded back in at the next Message Analyzer startup.

Message Analyzer keeps track of the windows that you display and the positions where you dock them in a configuration file that exists in the following location:

```
%LocalAppData%\Microsoft\MessageAnalyzer\app.WindowLayoutAsset.cfg
```

When you close the Message Analyzer application, the configuration file is updated so that you can resume your latest window layout on the next Message Analyzer restart. Note that this file as well as other configuration files in this location are now versioned in newer Message Analyzer releases.

See Also

Working With Message Analyzer Profiles

Working With Message Analyzer Profiles

71 minutes to read

To analyze message data that you load from saved files through a Data Retrieval Session, Message Analyzer enables you to choose different data viewer and view **Layout** configurations that provide various analysis contexts based on the viewing formats in which you present data. These configurations facilitate different perspectives on message data to enhance your analysis process. Because you have many options when selecting different data viewers and layouts, determining which is the most appropriate for the data you are examining could be challenging. Given that the viewing components you select can expose the data in different ways, it is important that you make the most appropriate choice to maximize your analysis capabilities. But if your experience with Message Analyzer is limited, your success in this effort will largely depend on trial-and-error methods.

For this reason, Message Analyzer now provides the **Profiles** feature, which enables you to utilize a set of built-in **Profiles** that contain specific viewer and layout presets that activate whenever you are loading data from specific types of input files. The **Profiles** feature also enables you to configure your own custom-designed **Profiles** so that you have the option to specify the viewers and layouts in which you want to expose your data. When you are configuring a **Profile**, you can associate a supported input file type with the **Profile** by making a selection from a drop-down list. After you save your **Profile**, it automatically applies the specified viewer/layout configuration to your Data Retrieval Session whenever you load data from the specified file type. Generally, the manner in which a custom **Profile** or built-in **Profile** is applied by Message Analyzer is functionally identical. The notable difference between them is that built-in **Profile** configurations are ReadOnly and predefined by Microsoft, while all custom **Profiles** are editable and predefined by you.

A simple scenario in which you could use a built-in **Profile** might be if you regularly analyze event trace log (ETL) files for specific types of information that require a particular view of data that quickly exposes the information you need to examine for ETW analysis. To display a typical view configuration, Message Analyzer enables you to use the built-in **Network Monitor Profile** for *.etl files, which defines a data viewer and layout configuration that is suitable for analysis of ETL data. When this **Profile** is enabled and you load data from an ETL file, Message Analyzer will automatically present the viewing configuration described ahead in "Exploring the Configuration of a Built-In Profile". If you do not want this **Profile** to activate when you are loading data from an ETL file, you can simply disable it, as described in Enabling and Disabling Profiles; or you can create a new **Profile** that specifies your own configuration, as described in Configuring a New Profile.

> **TIP**
>
> Network Monitor users who are new to Message Analyzer can create a familiar analysis environment by employing one of several **Network Monitor** built-in **Profiles** that are available. For example, when the **Network Monitor Profile** for a *.cap or *.etl file type is enabled, the resulting viewer and layout configuration provides the look and feel of a Capture or ETL file opened in Network Monitor. However, the viewer and layout configuration of these and other **Network Monitor Profiles** is suitable for any Message Analyzer user who wants to analyze such data.

Exploring the Configuration of a Built-In Profile Although all Microsoft-defined **Profiles** are ReadOnly, you can still explore the viewer and layout configuration of any ReadOnly **Profile** by first selecting it in the **Advanced Profiles** list and then clicking the **Edit Profile** button to display the **Profile** configuration. For example, you could explore any of the **Performance Top Down Profile** configurations, so that you can better understand the internal workings. You may find it helpful to

review the built-in **Profile** configurations when you create your own **Profile**, because you can use an existing **Profile** to create a baseline configuration that you can customize.

To explore a specific **Profile** such as the **Network Monitor Profile** for ETLs, select it in the **Advanced Profiles** list and then click **Edit Profile** and you will see the following configuration of viewers and layouts for this particular **Profile**. Note that the viewers described in the list that follows are common to all **Profiles**, built-in or custom-designed; however, the basic analysis context that is provided by each of these viewers is significantly enhanced by the view **Layouts** that are applied to them.

- **Analysis Grid** viewer — is set to use the **Network Monitor Viewpoints Layout**, a description for which is provided in Applying and Managing Analysis Grid Viewer Layouts. Because the **Analysis Grid** viewer is set as the default viewer for this **Profile**, the **Analysis Grid** viewer should automatically display with this particular layout shortly after you load data from an ETL file.
- **Grouping** viewer — is set to use the **Network Address and Ports Layout**, a description for which is provided in Understanding the Built-In Grouping View Layouts. Because the **Automatically open Grouping Viewer** check box is selected in this **Profile**, the **Grouping** viewer should automatically display with this layout shortly after you load data from an ETL file.
- **Chart** viewer — is set to use the **Top TCP/UDP Conversations by Message Count** view **Layout**, a description for which is included in the Chart Viewer Layouts section of this Operating Guide. Because **Charts** do not display by default in any of the built-in **Profiles** , you will need to manually launch them by selecting the **Default** item from the **Charts** drop-down list in the **New Viewers** drop-down list on the global Message Analyzer toolbar. As a result of this action, the **Layout** that is specified in the **Charts Layout** section of this **Profile** conveniently displays in a separate session tab.

By default, a number of built-in **Profiles** that Message Analyzer provides are enabled, which means that whenever you load data from a file type that is associated with one of these **Profiles**, for example, a .cap, .etl, or .log file, the **Profile** will automatically activate and present the data viewer and layout arrangement that it is configured to provide to your Data Retrieval Session results. If none of the built-in **Profiles** meet your requirements, you have the option to create one or more of your own by using the **Add Profile** feature on the toolbar under the **Advanced Profiles** section of the **Profiles** tab in the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. If the file type from which you are loading data into Message Analyzer is not associated with any existing and enabled **Profile**, then viewer and layout configurations are not impacted by the **Profiles** feature.

What You Will Learn In the remaining topics of this section, you will learn more about understanding and working with **Profiles**:

Understanding the Built-In Profiles
Applying and Managing Profiles
Enabling and Disabling Profiles
Configuring a New Profile
Example of Configuring a Profile to Create a Targeted Analysis Environment
Editing and Removing Profiles

Understanding the Built-In Profiles

The table that follows describes built-in **Profiles** that are provided by Message Analyzer along with

the associated input file types that activate the application of such **Profile** configurations when you are loading data into Message Analyzer. On the **Profiles** tab of the **Options** dialog that is shown in Using Message Analyzer Profiles, you may notice that there are several **Profiles** that have the same name but apply to different **File Types**. In some cases, the viewer and layout configurations are identical, while others vary. In cases where the configurations are identical in several **Profiles**, repetitive descriptions of these are omitted in the table that follows.

> **IMPORTANT**
> If you enable more than one **Profile** that applies to the same **File Type**, for example, capture (*.cap) files, Message Analyzer determines which **Profile** is applied to your loaded trace results based on an internal algorithm.

Table 14. Message Analyzer Built-In Profiles

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
Performance Top Down	.cap	Enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **Performance Top Down** view **Layout** populated with data, whenever you load data into Message Analyzer from a capture (*.cap) file for performance analysis. Also displays the **Grouping** viewer with the **Process Name and Conversations** view **Layout**, given that the **Automatically open Grouping Viewer** check box is selected in this **Profile**. To display the **Chart** configuration specified in this **Profile**, you will need to manually highlight the **Chart** item in the **New Viewer** drop-down list on the global Message Analyzer toolbar and then select the **Default** item in the list. This action will display the **Top TCP/UDP Conversations By Message Count** view **Layout** for the **Chart** viewer in a new session viewer tab. This latter **Layout** uses a **Table** grid visualizer component. **Usage Overview** — the advantages of the viewer and layout configuration of this **Profile** are that first, it provides the **Analysis Grid** viewer as a standard environment for detailed analysis that includes quick access to diagnosis errors and top-level messages that encapsulate message stacks, fragments, and any Operations. It also modifies the default **Layout** for this viewer, so that the configuration includes the **TopModule**, **TopSummary**, and **ResponseTime** columns, and also sorts the **TimeElapsed**

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		column in descending sort order. This sorted configuration can highlight performance issues by exposing messages with the highest elapsed time, possibly indicating delays in receiving message fragments. In addition, because **ResponseTime** data is available, you can correlate the time of the first server response to a request message with **TimeElapsed** data, to help determine whether performance issues are related to a server or to network latency. **More Information** **To learn more** about the **ResponseTime** annotation for Operations, see Average Elapsed Time for Operations. With the **Grouping** viewer **Process Name and Conversations Layout**, you can obtain a view that is similar to the Network Monitor Conversation Tree, in that the groupings enable you to isolate messages based on process name, process ID, network, and transport. With the **Top TCP/UDP Conversations By Message Count** view **Layout** for the **Chart** viewer, you can quickly obtain a summary of the Transport Layer conversations that carried the highest message count from data in the **Count** column, which is sorted in descending order. Other statistics are also included in this **Layout** such as **Bytes**, which indicates the total payload byte volume of all messages (containing this property) that are associated with each conversation; the data transmission rates in bytes-per-second (BPS) and kilobytes-per-second (KBs); along with data columns for conversation **StartTime**, **EndTime**, and **Duration**. After you display the **Layout** for this **Chart**, you might redock the **Chart** session tab, as described in Redocking Data Viewers and Tool Windows, so that it appears next to the **Analysis Grid** viewer. You can then drive selection of messages in the **Analysis Grid** viewer by double-clicking conversation data lines in the **Top TCP/UDP Conversations By**

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		Message Count view **Layout**. You can also select different groups in the **Grouping** viewer to correlate group messages with the results that display in both this **Chart** viewer and the **Analysis Grid** viewer. A summary of the type of information you can derive or the analysis that you can perform with the viewer and layout configuration of this **Profile** includes the following: • Data that can help you pinpoint performance problems related to slowly responding servers and/or network latency through analysis of **ResponseTime** and **TimeElapsed** data in the **Analysis Grid** viewer. • Diagnostic information that you can obtain in the **DiagnosisTypes** column of the **Analysis Grid** viewer. Diagnosis messages can quickly expose problem areas and guide the direction in which your inquiries should proceed for the resolution of issues. • Processes that consume high bandwidth, which can indicate potential problem areas, as exposed in the **Grouping** viewer by a summary view of network conversation message volumes and the Transport Layer ports over which they were carried for each **ProcessId** within **ProcessName** groups. • Network conversations with the highest traffic volumes and the ability to drill down into nested subgroups to view lower-level data in the grouped configuration. Also enables correlation of **Grouping** viewer data with **Analysis Grid** viewer and **Chart** messages through group selection. • Transport Layer performance statistics in the **Top TCP/UDP Conversations By Message Count** view **Layout** for the **Chart** viewer that exposes the top conversation message

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		volumes, data transmission rate, payload levels, and conversation durations. These statistics can pinpoint the conversations — and hence the computers — that may be experiencing performance problems. **Analysis Example** — from the data that you can obtain with the viewing and layout configuration provided by this **Profile**, you might be able to isolate the following types of issues: ◦ High message volumes from TCP Retransmits could be an indication of dropped packets, possibly caused by computer firewall rules, the network, or some other TCP issue. You can determine whether this may be the case by viewing diagnosis message descriptions for TCP messages of interest. ◦ High payload levels might point to a Windows scaling issue that results in dropped packets. ◦ Low data transmission rates at the client computer could be an indication of network delays. ◦ High conversation durations might combine several of these factors.
Network Monitor	.cap	Enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **Network Monitor** view **Layout** populated with data, whenever you load data into Message Analyzer from a *.cap file for process and performance analysis. Also displays the **Grouping** viewer with the **Process Name and Conversations** view **Layout**, given that the **Automatically open Grouping**

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		Viewer check box is selected in this **Profile**. To display the **Chart** configuration specified in this **Profile**, you will need to manually highlight the **Chart** item in the **New Viewer** drop-down list on the global Message Analyzer toolbar and then select the **Default** item in the list. This action will display the **Top TCP/UDP Conversations by Message Count** view **Layout** for the **Chart** viewer in a new session viewer tab. This latter **Layout** uses a **Table** grid visualizer component. **Usage Overview** — the advantages of the viewer and layout configuration of this **Profile** are that first, it provides the **Analysis Grid** viewer as a standard environment for detailed analysis that includes quick access to diagnosis errors and top-level messages that encapsulate message stacks, fragments, and any Operations. It also simulates the default Network Monitor view by including data columns such as **TimeDelta** (**TimeOffset** in Network Monitor) to indicate message running times, **ProcessName**, **Source** and **Destination** IP addresses, **Module**, and **Summary**. Note that Message Analyzer now captures process name data natively with any ETW provider, so that you can view process information with this **Layout** from any .cap trace file that contains it. With the **Grouping** viewer **Process Name and Conversations Layout** you can obtain a view that is similar to the Network Monitor Conversation Tree, in that the groupings enable you to isolate messages based on process name, process ID, network, and transport. With this configuration, you can view network conversations and the ports over which they were carried for each process ID within a process name group. With the **Top TCP/UDP Conversations By Message Count** view **Layout**, you can quickly obtain a summary of the Transport Layer conversations that carried the highest message count from the data displayed in the **Count** column, which is sorted in

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		descending order. Other statistics are also included in this **Chart** such as **Bytes**, which indicates the total payload byte volume of all messages containing this property that are associated with each conversation; the data transmission rates in bytes-per-second (BPS) and kilobytes-per-second (KBs); along with data columns for conversation **StartTime**, **EndTime**, and **Duration**. After you display this **Layout**, you might redock the **Chart** session tab, as previously described, so that you can observe the interactive display of data in multiple viewers based on message selection. **Analysis Example** — a sample of the type of information you can derive from the viewer and layout configuration of this **Profile** for analysis includes the following: • Process name information in the **Analysis Grid** viewer for which you can analyze the **Details** of an associated message. In addition, you can correlate such a message to a **Grouping** viewer group with the use of the **Find in Grouping Viewer** command located on the **Analysis Grid** viewer toolbar. For example, by selecting this command for a particular message of interest, you can highlight the **Network** or **Transport** group with which a **ProcessId** group is associated in the **Grouping** viewer, to expose the computer conversation and message volumes associated with the process name that displayed in the **Analysis Grid** viewer. • Diagnosis errors that can provide additional insights into the direction in which your analysis should proceed. You can quickly view these errors in the **DiagnosisTypes** column of the **Analysis Grid** viewer by sorting this column in descending order, as indicated by the down arrow near the column

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		label. You can also obtain a summary view of diagnosis message types, error descriptions, modules, and associated message counts across a set of trace results from the **Diagnostics Tool Window**. • A group-filtered view of messages and volumes associated with each of the following: ◦ Process name ◦ Process ID ◦ IP conversation ◦ Transport for messages in each IP conversation • Network conversations with the highest traffic volumes for a particular process, which could be a flag for further investigation. • Correlation of data in any **Grouping** viewer group with **Analysis Grid** viewer and **Chart** viewer messages through group selection. By manipulating the way data displays, you can achieve a unique enhancement to your analysis perspective. • TCP performance statistics from the **Top TCP/UDP Conversations By Message Count** view **Layout** for the **Chart** viewer, which includes top conversation message volumes, data transmission rate, payload levels, and conversation durations across a set of trace results, as previously described in the Analysis Example section of the **Performance Top Down Profile**. • Field data **Details** for any selected message, including information at the **Capfile** layer.
Network Monitor	.etl	Enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **Network Monitor** view **Layout** whenever you load data into Message Analyzer from a *.etl file for event log analysis. Also displays the **Grouping** viewer with the **Network Address and Ports** view

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		Layout, given that the **Automatically open Grouping Viewer** check box is selected in this **Profile**. To display the **Chart** configuration specified in this **Profile**, you will need to manually highlight the **Chart** item in the **New Viewer** drop-down list on the global Message Analyzer toolbar and then select the **Default** item in the list. This action will display the **Top TCP/UDP Conversations By Message Count** view **Layout** for the **Chart** viewer in a new session viewer tab. This latter **Layout** uses a **Table** grid visualizer component. **Usage Overview** — the advantages of the viewer and layout configuration of this **Profile** are that first, it provides the **Analysis Grid** viewer as a standard environment for detailed analysis that includes quick access to diagnosis errors and top-level messages that encapsulate message stacks, fragments, and any Operations. It also simulates the default Network Monitor view layout, which provides **ProcessName** information, as previously described in the **Network Monitor Profile** for .cap files. Also, with the **Grouping** viewer **Network and Addresses** view **Layout**, you can correlate IP conversations and the TCP/UDP port numbers that carried them, as presented in the **Network** and **Transport** groups. With the **Top TCP/UDP Conversations By Message Count** view **Layout** for the **Chart** viewer, you have access to TCP performance statistics that provide data to help you to isolate performance issues, as described earlier in the **Performance Top Down Profile** for .cap files. **Analysis Example** — a sample of the type of information you can derive or the analysis you can perform with the viewer and layout configuration of this **Profile** is specified earlier in the **Performance Top Down Profile** for .cap files, with exception of **ProcessName** and **ProcessId** information, which is not available in the **Grouping** viewer **Layout** for this **Profile**.

		this **Profile**.
PROFILE NAME **File Sharing SMB**	**APPLICABLE FILE EXTENSION** .cap	**DESCRIPTION** Enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **SMB Flat** view **Layout** whenever you load data into Message Analyzer from a *.cap file for SMB analysis. Also specifies the **Grouping** viewer with the **File Sharing SMB/SMB2** view **Layout**; however, you will need to manually select the **Grouping** viewer from the **New Viewer** drop-down list on the global Message Analyzer toolbar and then select the **Default** item in the list to display this viewer with the specified **Layout**, given that this **Profile** is not configured to automatically display the **Grouping** viewer. To display the **Chart** viewer that is configured in this **Profile**, you will need to manually highlight the **Chart** item in the **New Viewer** drop-down list on the global Message Analyzer toolbar and then select the **Default** item in the list. This action will display the **SMB Top Talkers** view **Layout** for the **Chart** viewer in a new session viewer tab. This latter **Layout** uses a **Table** grid visualizer component. **Usage Overview** — the advantages of the viewer and layout configuration of this **Profile** are that first, it provides the **Analysis Grid** viewer as a standard environment for detailed analysis that includes quick access to diagnosis errors and top-level messages that encapsulate message stacks, fragments, and any Operations. In addition, it includes data columns such as **TimeDelta**, **Source**, **Destination**, and **Summary** information while exposing several SMB data fields that you can quickly reference for further analysis of message details. Data columns that are significant for SMB analysis are the **SessionIdName**, **TreeIdNameReference**, **FileNameReference**, and **Header.MessageId** columns. They provide the following information: - **SessionId** — provides a value that uniquely identifies each session that is multiplexed over a single SMB connection. - **TreeId** — provides a value that uniquely identifies a connection between a Common Internet File

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		System (CIFS) client and a share on a remote CIFS server. - **FileNameReference** — provides the name of the file resource/s upon which SMB operations were performed. - **MessageId** — provides a value that uniquely identifies an SMB request and response pair among all messages that are sent across a common SMB connection. The **Grouping** viewer enables you to view the message volume per session, as distinguished by a **SessionIdName** group, among potentially multiple sessions over a single SMB connection. Drilling down further, you can view specific share connections (TreeIds) via the nested **TreeIdName** groups along with the nested **FileName** groups under each parent **TreeIdName** group. At each group level, the **Grouping** viewer enables you to examine the traffic volumes associated with each group in the nested configuration and to interactively drive display of messages associated with any selected group into the **Analysis Grid** viewer for further investigation of message details. The **SMB Top Talkers** view **Layout** for the **Chart** viewer enables you to examine a summary of IP conversations (via address pair sets) sorted by message count from highest to lowest, along with other statistics that include **Bytes**, which indicates the total payload byte volume of all messages containing this property that are associated with each conversation; the data transmission rates in bytes-per-second (**BPS**) and kilobytes-per-second (**KBs**); along with data columns for conversation **StartTime**, **EndTime**, and **Duration**. **Analysis Example** — for instance, if SMB write or read operations are taking a long time, possibly indicated by a high **Duration** value (sort this column in descending order for the best view), you may be able to isolate a poorly performing computer where this is occurring by observing the session duration, message count, and/or data transmission rate that is associated with the conversation in which such a computer is engaged.

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		You can also interactively drive the display of data in the **SMB Top Talkers** view **Layout** for the **Chart** viewer and the **Analysis Grid** viewer, from any group that you select in the **Grouping** viewer, for further correlation of data, as described earlier. For best interactive results, redock the **Chart** session viewer tab next to the **Analysis Grid** viewer. You might also keep in mind that errors may be occurring, which you can view in the **DiagnosisTypes** column of the **Analysis Grid** viewer, as described earlier in the **Network Monitor Profile** for .cap files in this table.
File Sharing SMB Perf	.cap	Enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **File Sharing Perf SMB2/SMB** view **Layout** whenever you load data into Message Analyzer from a *.cap file for SMB analysis. Also specifies the **Grouping** viewer with the **File Sharing SMB/SMB2** view **Layout**; however, you will need to manually select the **Grouping** viewer from the **New Viewer** drop-down list on the global Message Analyzer toolbar and then select the **Default** item in the list to display this viewer with the specified **Layout**, given that this **Profile** is not configured to automatically display the **Grouping** viewer. To display the **Chart** viewer that is configured in this **Profile**, you will need to manually highlight the **Chart** item in the **New Viewer** drop-down list on the global Message Analyzer toolbar and then select the **Default** item in the list. This action will display the **SMB Service Performance** view **Layout** for the **Chart** viewer in a new session viewer tab. This latter **Layout** uses a **Table** grid visualizer component. **Usage Overview** — similar to the **File Sharing SMB Profile**, the advantages of the viewer and layout configuration of the **File Sharing SMB Perf Profile** are that first, it provides the **Analysis Grid** viewer as a standard environment for detailed analysis that includes quick access to diagnosis errors and top-level messages that encapsulate

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		message stacks, fragments, and any Operations. It also includes the same data columns as the **File Sharing SMB Profile** for .cap files, with the exception of the **TimeDelta** column, which is replaced with the **ResponseTime** column in this **Profile**. Therefore you can obtain similar values and statistics with both of these **Profiles**, although with the **File Sharing SMB Perf Profile**, you can also assess the server response times to SMB2 request messages, as they are conveniently located in the **ResponseTime** column of the **Analysis Grid** viewer that you can add with **Field Chooser**. **Analysis Example** — if you correlate **ResponseTime** and **TimeElapsed** data in the **Analysis Grid** viewer, you can determine whether performance is being compromised by a slowly responding server or by network latency, as described earlier in the **Performance Top Down Profile**. Other key data fields for the viewers in this **Profile** consist of **SessionId**, **TreeIdReference**, and **FileNameReference**, which are provided in **Analysis Grid** viewer as columns, and in the **Grouping** viewer as equivalent groups. See the **File Sharing SMB Profile** for more information about these fields. Also keep in mind that diagnosis messages may be helpful in determining the cause of performance issues.
Event Log	.evtx	Enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **Event Log** view **Layout** whenever you load data into Message Analyzer from a *.evtx file for event analysis. You can also display the **Grouping** viewer with the **Event Viewer** view **Layout**; however, because the **Automatically open Grouping Viewer** check box is unselected in this **Profile**, you will need to manually select the **Grouping** item from the **New Viewer** drop-down list on the global Message Analyzer toolbar to display the indicated configuration. To display the **Event Log IDs** view **Layout** for the **Chart** viewer, manually select the **Chart** item in the **New Viewer** drop-down list to display the indicated **Layout**. This latter **Layout** uses a **Bar** element visualizer component.

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		Usage Overview — the advantages of the viewer and layout configuration of this **Profile** are that first, it provides a basic analysis environment with the **Analysis Grid** viewer for viewing event data. The **Event Log** view **Layout** enables quick access to standard event information. Much of this information is declared in an Event Descriptor, which in turn is typically defined by an ETW provider manifest, as described in the ETW Framework Conceptual Tutorial. The event information that can populate this **Analysis Grid** viewer **Layout** can include **EventID**, **Version**, **Channel** (target audience), error **Level**, and **Opcode**. Keyword values are also typically a part of event definitions and usually reside in an event manifest. Other important information that is exposed by this **Layout** includes the **ProcessId**, ETW **ProviderName**, and the actual **EventData** that tells you the current state of an application or some process. **Note:** You can view Keywords for any *.evtx log in the **Details** window. You can also add a **Keywords** column to the **Analysis Grid** viewer by right-clicking the **Keywords** field in the **Details** window and then selecting **Add 'Keywords' as Column** in the context menu that displays, that is, after initially selecting an event/message in the **Analysis Grid** that defines Keywords. Also, the **Grouping** viewer with the **Event Viewer Layout** contains the following four groups in a nested configuration for every data set that is defined by a unique top-level field value: - **ProviderName** — this top-level field is the name of the ETW provider that raised events and wrote them to the ETW session from which your data is displaying. - **Level** — this field can include error Levels in the range of 1-5, for example, Critical (1), Error (2), Warning (3), and so on. - **Channel** — this field displays the target audience for the event/s and is specified in an ETW provider manifest. - **EventID** — this field specifies the ID for events that were written by an ETW provider.

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		As different **ProviderName** values are detected in the trace results, additional grouped data sets are created and organized by Message Analyzer to expose the different values in the above specified nested group configuration. When you select a group node in the **Grouping** viewer for any data set, the messages that correspond with that group node are filtered to the **Analysis Grid** viewer, so that you can analyze all the messages associated with a common group value. For example, this could be a specific error **Level**, **Channel**, or **EventID** value. This provides a unique way of organizing the trace data into summary groups that enable you to interactively correlate different aspects of your data with the analysis context of data displayed in the other viewers that are configured by this **Profile**. The **Event Log ID** view **Layout** for the **Chart** viewer enables you to view the message volume — ordered from the highest to the lowest volume — that is associated with **EventIDs** that were found in the .evtx log. The data is displayed in a **Bar** element visualizer component that provides an at-a-glance view of the relative distribution of message volume per **EventID** across a set of trace results. This enables you to make a quick visual assessment of which events involved the highest message count, which could be a flag for further investigation. As you click any bar element in this **Layout**, messages associated with that element are highlighted in the **Analysis Grid** viewer. This same result occurs if you select **EventID** groups in the **Grouping** viewer, provided that the **Grouping** viewer is in **Selection Mode**. Otherwise, when the **Grouping** viewer is in **Filtering Mode**, messages associated with the clicked group will be filtered to the **Analysis Grid** viewer for further examination and to the **Chart** viewer as well. **Analysis Example** — from the viewer and layout configuration of the **Event Log Profile**, you can derive the following types of information which can be

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		significant to the analysis process: - The message volumes associated with the events of a particular message provider, as exposed in the group configuration of the **Grouping** viewer. Message volumes per **EventID** are also exposed in the **Chart** viewer for this **Profile**. High volumes might point to an overburdened system component or application that is issuing a lot of event traffic or experiencing a high rate of errors. Sparse traffic might be an indication of dropped packets due to misconfigured ETW Session buffer settings, as described in Specifying Advanced ETW Session Configuration Settings. - The ETW provider that is writing the events, as exposed by the **ProviderName** field. This can identify the message provider for a particular component, application, or subsystem that may be experiencing performance issues. - The error levels and descriptions associated with each provider's messages, as exposed by the **Level** or **LevelDisplayName** and **Summary** fields, respectively. This can expose the severity of event errors, which can be a flag to examine any diagnosis messages that are associated with such errors. In turn, diagnosis message descriptions may expose an underlying issue. - The process ID associated with each event, as exposed by the **ProcessId** field, which could pinpoint a particular application or process that is experiencing errors, erratic behavior, or sluggish performance. - The event Keywords configured in the event manifest for the ETW provider, as exposed by the **Keywords** field. Only the Keywords that were specified in the ETW provider manifest are reported to the ETW Session and subsequently recorded in the .evtx log, that is, if such events were written by the ETW provider in response to some error condition or state of an application or system component. Such information can highlight a problem area for further investigation. - The Diagnosis messages associated with each event, as exposed in the **DiagnosisTypes**

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		and **Summary** columns of the **Analysis Grid** viewer. Diagnosis messages consist of four types, as described in the Diagnosis Category topic. For example, a Diagnosis message might indicate that a particular event could not be parsed by Message Analyzer due to invalid data (a Parsing error type) or that an event does not align with its manifest definition (a Validation error type). **Tip:** To enhance the interactive analysis context for the viewers and layouts of this **Profile**, you might redock the **Chart** viewer tab alongside the **Analysis Grid** viewer so you can observe the interaction between **Grouping** viewer group node selection and the display of messages in the **Analysis Grid** and the **Chart** viewer, as described in the first item of this table.
Fiddler Traces	.saz	Enable this **Profile** to display the **Analysis Grid** viewer as the default viewer along with the **Fiddler SAZ** view **Layout** whenever you load data into Message Analyzer from a Fiddler *.saz file for HTTP analysis. Note that the data exposed in this viewing configuration closely resembles the Fiddler Web Debugger analysis environment, although field names are different. You can also display the **Grouping** viewer with the **Fiddler Grouping** view **Layout**; however, because the **Automatically open Grouping Viewer** check box is unselected in this **Profile**, you will need to manually select the **Grouping** item from the **New Viewer** drop-down list and then select the **Default** item in this list, as previously described, to display the indicated configuration. Likewise, to display the **HTTP Content Type Volumes** view **Layout** for the **Chart** viewer, manually select the **Chart** item in the **New Viewer** drop-down list and then select the **Default** item in this list to display the indicated **Layout**. This **Layout** uses a **Bar** element visualizer component. **Usage Overview** — the main advantage of the viewer and layout configuration of this **Profile** is that it provides the **Fiddler SAZ** view **Layout** in the **Analysis Grid** for in-depth analysis of HTTP

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		messages from Fiddler traces in a simulated Fiddler debugging environment. Some of the most significant information that is exposed in this **Analysis Grid** viewer **Layout** for analysis consists of the following: - HTTP response code, as exposed in the **StatusCode** column. - HTTP verbs, as exposed in the **Method** column. - Uniform resource identifier (URI) information such as the host, absolute path to resources, and URIs, as exposed in the **uri.host**, **uri.abspath**, and **uri columns**, respectively. - Packet length, equal to the header + payload in bytes, as exposed in the **PayloadLength** column. - Content caching directives, as exposed in the **Headers.cache-control** column. - Content type, process name and ID, and payload value information, as exposed in the **ContentType**, **SessionFlags.x-processinfo**, and **Payload** columns, respectively. The **Grouping** viewer isolates some of this same information from a Fiddler trace into groups, where you can view the message volume that is associated with each top-level process name and ID group (**SessionFlags.x-ProcessInfo**), the hosts that handled each request as indicated in the nested **Uri.Host** group under a particular process name and ID group, along with the number of messages associated with each host group. The **HTTP Content Type Volumes** view **Layout** for the **Chart** viewer provides a visual indication of the relative volumes of HTTP content type payload lengths in bytes for each content type, along with the relative distribution of volume for each content type, from the highest to lowest values. This enables you to see at a glance which byte volumes are the largest for any particular content type. In turn, this can provide an indication of the loads being carried by responding web servers. **Analysis Example** — an example of how you might use these viewer

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		and layout configurations as tools for analysis is to first sort the **PayloadLength** column of the **Analysis Grid** viewer in descending order so that you can see which messages had the highest packet length. You can then correlate that information with the following: - The HTTP request type, as specified in the **Method** column. - Process name and process ID, as specified in the **SessionFlags.x-ProcessInfo** column. - Content type associated with messages of interest, as specified in the **ContentType** column. - The web server host and specific resources that were requested by a client, as specified in the **Uri.Host** and **Uri.AbsPath** columns. - Status of HTTP response messages, as specified in the **StatusCode** column. In summary, the information that you obtain from the **Analysis Grid** viewer with this correlation can expose the types of request messages that are associated with a particular process, the specific type of content involved, the hosts from which resources were retrieved, along with the success of the operations. With this data, you may be able to pinpoint a web server that is under stress, potentially from servicing a high volume of client requests for a particular content type. Moreover, you can use the **Find in Grouping Viewer** command in the **Analysis Grid** viewer's right-click context menu for particular messages with various **PayloadLength** values, so that you can locate them in the **Grouping** viewer for a quick correlation of associated process and host data. In addition, you can interactively and simultaneously drive the display of messages in the **Analysis Grid** viewer and the **HTTP Content Type Volumes** view **Layout** for the **Chart** viewer by selecting **ProcessInfo** group nodes in the **Grouping** viewer that contain various message volumes, providing that the Grouping viewer is in the **Filtering Mode**. This enables you to isolate the associated group messages in these other viewers to take

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		advantage of their analysis capabilities. Note that if the **Grouping** viewer is in **Selection Mode**, the messages that are associated with a selected **ProcessInfo** group will be highlighted in the **Analysis Grid** viewer only. Alternately, you can double-click a bar element of a certain content volume in the **HTTP Content Type Volumes** view **Layout** for the **Chart** viewer to isolate the messages represented in that bar element to a separate instance of the **Analysis Grid** viewer for review of message **Details**. You can also use the **Find in Grouping Viewer** command on the isolated **Analysis Grid** messages to expose and correlate the process information in the **Grouping** viewer with the hosts involved. Set the **Grouping** viewer to the **Selection Mode** for this operation to work the best. By examining this information in the indicated ways, you may be able to determine that one or more responding web servers are carrying large loads, which could expose performance issues that include sluggish response times.
Text log files	Common file extension:	**Important:** Because Message Analyzer has multiple built-in **Profiles** for different logs that are all associated with the same *.log* file type designation, you will need to open the log file types described below in one of the following ways, otherwise the correct view **Layout** for the **Chart** viewer will not display after you load data from these logs. By specifying a text log configuration file in the actions that follow, Message Analyzer can differentiate between the built-in **Profiles** for different Log files, so that the right **Profile** is applied: • Use the **New Session** dialog for a Data Retrieval Session, from where you can select the configuration file for the specific .log type that contains the data you are importing.--- **More Information To learn more** about using text log configuration files to open .log files, see Opening Text Log Files.--- • Use the **Open** command or **Recent Files** list on the

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		Message Analyzer **Start Page**, providing that you have correctly set the default text log configuration file on the **General** tab of the **Options** dialog for the .log type you are opening. If you do not have a default configuration file set on the **General** tab of the **Options** dialog and you attempt to open any .log file, for example, by using the **Open** command or **Recent Files** list on the **Start Page**, or by opening the log from Windows Explorer, Message Analyzer presents you with the input configuration for a Data Retrieval Session so that you can specify the **Text Log Configuration** file that is appropriate for the log file type that contains the data you are importing. If you are uncertain about which **Text Log Configuration File** to specify for any of the following logs, you can view the **Profile** configuration for each log by selecting an appropriate **Profile** on the **Profiles** tab of the **Options** dialog and then clicking the **Edit** button on the **Advanced Profiles** toolbar. **Note:** Message Analyzer parses the log files that are described in this section with the use of configuration files, which contain OPN code that is specifically designed to parse such files. When the OPN code in a configuration file for a particular log type is compiled, an OPN module for that log is added as a new node in the **Field Chooser Tool Window**. If you expand such a node, you will have access to other log fields that were parsed by the OPN code. You can then add any of these fields as a new data column in the **Analysis Grid** viewer by double-clicking it. The resulting data that is displayed can provide additional information to support your analysis process. **To learn more** about the **Field Chooser**, see the topic Field Chooser Tool Window.
IIS Logs	.log	Enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **IIS** view **Layout**, whenever you load data from an IIS .log file for analysis of client and server data in IIS logs.

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		You will need to manually open the other viewers that are configured in this **Profile** in the previously described manner, which includes the **Grouping** viewer with the **IIS** view **Layout**, and the **IIS Log HTTP Traffic Volume** view **Layout** for the **Chart** viewer. This latter **Layout** uses the **Bar** element visualizer component. **Note:** In the **IIS** view **Layout** of the **Analysis Grid** viewer, data fields that are associated with the client computer contain a "c" character in the prefix of the field name, while "cs" characters indicate a client-to-server transaction. Likewise, fields that are associated with the server contain an "s" character in the prefix of the field name, while "sc" characters indicate a server-to-client transaction, although you will only find fields with "sc" characters in the **Details Tool Window**. Examples from the **Analysis Grid** viewer **Layout** include **cs_method** and **s_port**, **Usage Overview** — the main advantage of the viewer and layout configuration of this **Profile** is that it provides data sets in several different interactive viewing configurations that expose the information you will need to analyze the logs of an IIS web server. Just from the **Analysis Grid** viewer alone with its **IIS** view **Layout**, you may be able to discover factors such as the following that might be contributing to web server stress: • Operations that took a long time to complete for a particular type of client request, as exposed in the **time_taken** column. **Tip:** Add this field as a column in the **Analysis Grid** viewer by right-clicking the **time_taken** field in the **Details** window and then selecting the **Add 'time_taken' as Column** command. This field exposes the length of time in milliseconds that an action took to complete. • High traffic volumes associated with any of the following: ◦ Sites that are being inundated with the

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		most client traffic, as exposed in the **s_sitename** column and correlated to **Grouping** viewer **s_port** (server port) message volumes. You can correlate a site name with message volumes on an input server port by clicking the **Find in Grouping Viewer** command on the **Analysis Grid** toolbar while a site name of interest is selected. This action should highlight an **s_port** group in the **Grouping** viewer and show the associated message volume. **Note:** The **Grouping** viewer must be displayed so that the **Find in Grouping Viewer** command is enabled. ◦ Target resources or services to be accessed, as exposed in the **cs_uri_stem** column. ◦ Queries/client requests, as indicated in the **cs_uri_query** column. ◦ Methods and operations to be performed, as exposed in the **cs_method** column. ◦ Specific users who are making requests, as exposed in the **cs_username** column. ◦ Potentially compromised client browsers or other applications that are sending erroneous or intermittent queries to the server, as exposed in the **csUser_Agent** column. At a minimum, the information

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		provided in this **Analysis Grid** viewer **Layout** could expose any of the following issues: - Slow server operations (high **time_taken** values) could be an indication that a high volume of requested operations are stressing web server resources, compromising performance, and reducing service availability. - High traffic volumes for specific sites can indicate that such sites are being overwhelmed by traffic and possibly dropping packets. - High traffic volumes associated with client queries, client methods, specific users, or target resources might be consuming web service availability time. The **Grouping** viewer exposes the client IP addresses that made the requests and the server ports that received the requests, along with the query message volume sent to the server by the client. If you also have the **IIS Log HTTP Traffic Volumes** view **Layout** for the **Chart** viewer displayed, you can view the relative distribution of traffic volume in bytes, from the highest to the lowest volume, for the server HTTP responses to each client query that the server received. The volume values in this **Layout** are based on the **sc_bytes** field for server responses, the values for which you can view in the **Details** window. This visualizer component provides a quick summary of the server response volumes in bytes that are associated with the queries requesting access to web server resources and services. Note that very high byte volumes could be a flag that points to the potential overload of one or more web servers. You can also use the **Grouping** viewer to interactively and simultaneously drive the display of messages in the **IIS Log HTTP Traffic Volumes** visualizer component and the **Analysis Grid** viewer, by group selection in the **Grouping** viewer. For example, if the **Grouping** viewer displays multiple **c_ip** groups of client addresses where requests were initiated, you can view the associated messages in the **Analysis Grid** viewer and

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		corresponding server response byte volumes in the **IIS Log HTTP Traffic Volumes** view **Layout** (to isolate the data for further analysis) by clicking those groups in the **Grouping** viewer. As previously described, you can also right-click any message in the **Analysis Grid** viewer and select the **Find in Grouping Viewer** context menu command to locate the group in the **Grouping** viewer with which a message of interest in the **Analysis Grid** viewer is associated. **Tip:** Additional IIS log fields are available for examination in the **Details** window, which includes server response data such as sc_status and sc_bytes.
Netlogon Logs	.log	Enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **Netlogon Log** view **Layout**, whenever you load data from a Netlogon .log file to analyze Netlogon data. You will need to manually open the other viewers that are configured in this **Profile** in the previously described manner, which includes the **Grouping** viewer with the **Netlogon Group by Message Type** view **Layout**, and the **Netlogon Message Types** view **Layout** for the **Charts** viewer. This latter **Layout** uses the **Pie** chart visualizer component. **Usage Overview** — the advantages of the viewer and layout configuration of this **Profile** consist of the following: - The **Analysis Grid** viewer with the **Netlogon Log Layout** provides summary data for each log file entry for a Netlogon .log file that includes message type information, along with data in other **Analysis Grid** fields that include **MessageNumber**, **Timestamp**, and **TimeDelta**. - The **Grouping** viewer with the **Netlogon Group by Message Type Layout** isolates messages into message type groups and provides the number of messages associated with each type. - The **Netlogon Message Types** view **Layout** for the **Chart** viewer provides a pie-slice visualizer that summarizes the relative percentage of message volumes for each message type in a Netlogon log file.

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		Analysis Example — in the **Summary** column of the **Analysis Grid** viewer, you will find a description that includes message type, error descriptions, and other descriptive data that is related to each message. As described many times in this table, you can associate any message in the **Analysis Grid** with groups in the **Grouping** viewer; for this **Profile** it would be the **msgType** group. Some of the important message types that are issued during the log on process and which you will typically find in a Netlogon log are as follows: • **MAILSLOT** — consists of an LDAP ping that enables a client to locate a domain controller with this type of message via RPC name pipes or with TCP as the transport. This message is in turn received by a logon server that has created a MAILSLOT file that the client message can write to, thus establishing client-to-server authentication communications for logon. Both client MAILSLOT and server response MAILSLOT messages are typically written to Netlogon logs so that you can view client and server communication records. • **DNS** — this type of message from a logon server can provide cache entries and annotations, or an indication of DNS status. • **CRITICAL** — this type of message from a logon server typically includes critical error information or status, such as an invalid domain name was pinged, a DNS query failed to return data, NetBIOS to IP address resolution failed, and so on. • **DIAGNOSIS** — this type of message from a logon server typically includes error information or status, such as DNS resolution failures, indications of authentication chain issues, setup problems, client queries that failed because the server could not service them, and failure to find a logon server. Note that

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		DIAGNOSIS messages can also be of type **CRITICAL**. • **SESSION** — this type of message provides a record of different logon session-level messages, such as the following: ◦ Messages associated with establishing a session, including domain controller discovery and machine password resets. ◦ Requests sent to the security account manager (SAM) and associated responses. ◦ Authentication success and failure messages. ◦ Messages requesting logon domain information. ◦ Queries for logon server capabilities. • **LOGON** — this type of message provides a record of successful logons, for example, Administrator and user logons, along with account and site names. Another type of **LOGON** message is the **WRONG PASSWORD** message, which can indicate the following: ◦ A password provided to a non-logon server, for example, a file server, was forwarded to an authenticating domain controller for validation. Also known as pass-through authentication. ◦ A password did not match the one held by an authenticating domain controller and was therefore passed to the primary domain controller for validation. • **MISC** — this type of message can include

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		Netlogon events being logged, user or machine account status, and other information requested by a client, such as the status of a logon server, domain controller name requests, and other information. It might also include server messages that advise a client of other sites with greater availability. • **PERF** — this type of message provides Netlogon performance counter information that includes data related to setting up a client-server session, the number of authentication timeouts that have occurred, and average semaphore hold times before authentication occurs. **Note:** Messages for various authentication types that can be detected by the Message Analyzer Netlogon text log parser include NTLM, Kerberos PAC, Digest, and so on. With the viewer and layout configuration of this **Profile**, you can very quickly isolate the above information during analysis to find problem areas. You can do this by clicking on each message type in the **msgtype** group of the **Grouping** viewer. From this action, you can effectively isolate the messages associated with each group in the **Analysis Grid** viewer, provided that the **Grouping** viewer is in **Filtering Mode**. If the **Grouping** viewer is in the **Selection Mode**, you can simply highlight the messages in the **Analysis Grid** viewer without introducing any filtering effects. If you also have the **Netlogon Message Types** view **Layout** for the **Chart** viewer displayed, you can click different **Pie** chart elements and drive the display of messages in the **Analysis Grid** viewer. When you do this, you can also achieve different interactions with the **Grouping** viewer depending on the mode it is in. These capabilities enable you to quickly zero-in on the specific data presented by different message types, which is very convenient

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		when you need to expose errors and other important information that is buried in a large log file. **More Information** **To learn more** about the Netlogon troubleshooting and the Netlogon parser that is used by Message Analyzer, see Diving into the Netlogon Parser (v3.5) for Message Analyzer on TechNet.
Cluster Logs	.log	Enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **Cluster Log** view **Layout**, whenever you load data from a Cluster .log file to expose fields that are key to analysis. You will need to manually open the other viewers that are configured in this **Profile** in the previously described manner, which includes the **Grouping** viewer with the **Cluster Logs** view **Layout**, and the **Cluster Levels** view **Layout** for the **Chart** viewer. This latter **Layout** uses the **Bar** element visualizer component. **Usage Overview** — the main advantage of the viewer and layout configuration of this **Profile** is that it provides several data sets in different interactive viewing configurations that expose information you will need to quickly isolate problem areas for further investigation of clustering issues. From the **Analysis Grid** viewer **Layout**, you can obtain an overview of key cluster log information through data columns such as **InfoLevel**, **Subcomponent**, **RemainingText**, **ProcessId**, and **ThreadId**. With this information, you can expose errors that may be occurring in a particular subcomponent of the Cluster Service, for example, the Failover Manager, Database Manager, Node Manager, or Global Update Manager; and you can also associate such errors with one or more **ProcessIds**. Sorting and grouping in the **Analysis Grid** viewer can organize the data in a way that speeds up analysis. For example, if you sort the **Subcomponent** column of the **Analysis Grid** viewer in ascending order, you can organize the log entries such that the entries for any particular component are

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		gathered together for easy viewing. Moreover, you can execute the **Group** command from the context menu that displays when you right-click the **Subcomponent** column header in the **Analysis Grid** viewer. The result of this operation provides a view of the data that encapsulates the message activity that occurred for various subcomponents of the Cluster Service into a separate "group" node that you can expand for further details. Likewise, if you **Group** the **InfoLevel** column, you will see a view of the data that encapsulates the message activity associated with the information level that exists for each log entry for debugging purposes. Note that a quick way to expose failures that might have occurred is to **Apply** a Filter such as `*RemainingText contains "failure"` from the Message Analyzer Filtering Toolbar that is located in the upper left sector of the **Analysis Grid** session tab . The results of this operation can point you to specific components where errors occurred, while also providing a description of what actually occurred. But probably the most useful way to display the data is with the default **Layout** of the **Grouping** viewer. This **Layout** enables you to isolate the different types of information levels that can be written by a Cluster Service component, which typically consist of informational (**INFO**), warning (**WARN**), error (**ERR**), and debug (**DBG**) levels. These informational levels are isolated by the top-level **InfoLevel** group in this **Layout**. The **Subcomponent** group is nested under the **InfoLevel** group and the **ProcessId** group is in turn nested under that. By organizing the data in this grouped configuration, this **Layout** enables you to very quickly assess all the information levels that occurred for each Cluster Service **Subcomponent** and the **ProcessIds** that are associated with the operations that were carried out. **Tip:** You can obtain a quick assessment of which information levels have the most log entry activity by opening up

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		the **Cluster Levels** view **Layout** for the **Chart** viewer. This data display provides an at-a-glance view of the relative distribution of message volume for each of the information types found in a Cluster log file. By double-clicking any bar element that represents a particular **InfoType**, you can display all the log entries that contain that type along with the Subcomponents with which they are associated. **Analysis Example** — the Global Update Manager (GUM) is a primary mechanism of the Clustering Service that keeps all cluster nodes up to date with the latest resource configurations stored in the Cluster database. It is also used by internal Cluster Service components, such as the Failover Manager (FM), Node Manager (NM), and Database Manager (DM), to replicate changes made to any node, which is usually initiated by a Cluster API call. The GUM is a heavy user of Cluster Service communication processes and is therefore a good starting point when troubleshooting clustering issues. To assess any issues that may have occurred with the GUM service, you can do the following: 1. Open the **Grouping** viewer and then display the **Cluster Logs Layout** in the previously described manner. 2. Click the **Collapse All** button on the **Grouping** viewer toolbar to display the top-level groups only, which in this case will be the data for **InfoLevel** groups that is derived from your Cluster log. 3. Click the expansion node of the **ERR** group to display the nested **Subcomponent** groups. 4. Scroll down to the **GUM** group and click it to display all the log entries that contain errors that were logged by the GUM service. If the **Grouping** viewer is in **Filtering Mode**, this action will filter and display the associated messages to the **Analysis Grid** viewer. If the **Grouping** viewer is in the **Selection Mode**, the same messages will simply be highlighted in the **Analysis Grid** viewer. 5. Observe the error descriptions under the **RemainingText** column

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		of the **Analysis Grid** viewer. For example, you might see that a GUM request resulted in an exception or other failure during the update process for a specific cluster node. 6. Obtain the **ProcessId** that is associated with any log entry that exposes an error, by right-clicking the log entry and then selecting the **Find in Grouping Viewer** command in the context menu that appears. The relevant process will be highlighted in the **ProcessId** group that is nested under the **Subcomponent** group. This information may provide some additional insights into which resources or other components were involved in the failed update process. 7. If no errors were logged in the **ERR** group for the GUM service, go to step 3 and perform these same operations for the **WARN** group. **Tip:** For hints of other potential problem areas, you can also review the **TimeDelta** column values for evidence of operations that took an exceptionally long time to complete.
Samba Logs	.log	If you are a developer who tests new Samba features or if you simply want to monitor Samba performance, you can enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **SysLog** view **Layout**, whenever you load data from a SambaSysLog .log file. You will need to manually open the other viewers that are configured in this **Profile** in the previously described manner, which includes the **Grouping** viewer with the **SysLog** view **Layout**, and the **SysLog Levels** view **Layout** for the **Chart** viewer. This latter **Layout** uses a **Bar** element visualizer component. **Usage Overview** — the main advantage of the viewer and layout configuration of this **Profile** is that it provides several data sets with varying analysis contexts that can quickly expose the Samba log entries where issues may be occurring. For example, in the **Analysis Grid** viewer, you can correlate the Samba debug levels with the Samba **functions** that wrote the log entries and which contain the **level** information, along with the Samba

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		source_file/s where the functions exist. A quick way to summarize this information might be to sort the **level** column so you can view all the log entries in a hierarchical manner according to **level** values. You can then correlate log entries that have the more critical **level** values with the associated **function** and **source_file** data. **Analysis Example** — you might consider taking advantage of the **Analysis Grid** viewer **Group** command to organize the data into separate hierarchical groups that each contain log entries with a common **level** value, so that you can evaluate the data in the context of identical **level** value groups. You can also nest additional groups under the **level** group, for example, a **function** group at the first nested level and a **source_file** group at the second nested level. Then, by drilling down to the **source_file** group you can expose the log entries that have been isolated according to the grouped configuration. To execute an **Analysis Grid** viewer **Group** command, right-click the header of a column such as **level** and select the **Group** item that appears in the context menu that displays. After you create a multiple group configuration in this manner, you can drag any group into a new position in the hierarchy to recast the data according to the new group organization that is created, so that you can obtain an alternate analysis perspective on the data. The **Grouping** viewer provides a similar grouping configuration; however, it also enables you to interactively drive selection of log entries in the **Analysis Grid** viewer based on group selection in the **Grouping** viewer. As previously described, if the **Grouping** viewer is in the **Selection Mode**, group selection will cause **Analysis Grid** viewer log entries to be highlighted; if the **Grouping** viewer is in the **Filtering Mode**, group selection will cause a filtered view of the log entries where all other entries are temporarily removed from the **Analysis Grid**, that is, until you click the **Reset** button on the **Grouping** viewer toolbar.

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		The advantage of the Grouping viewer is that you can isolate the log entry data to the top group, which is the Samba debug **level**, to the Samba **function** that wrote the log entry to the first nested group, and to the Samba **source_file** that contains the function in the last nested group. This grouped configuration enables you to prioritize your investigation based on the **level** values, which is a good starting point from where you can determine, in a hierarchical manner, the functions and source code that is associated with the most critical levels. SambaSysLog levels typically consist of the following: - 0 — Error - 1 — Warning - 2 — Notice - 3 — Information - 4 and above — Debug You might proceed by first clicking the **Collapse All** button on the **Grouping** viewer toolbar so you can immediately see all the debug levels that exist in the entries of your SambaSys log. Then click a **level** expansion node that is designated with a value such as '0' or '1' to expose the data for the underlying **function** and **source_file** groups. Next, make sure the **Grouping** viewer is in **Filtering Mode** by clicking the **Filtering Mode** icon on the **Grouping** viewer toolbar and then select a **function** group value of interest. The log entries associated with the selected function are filtered to the **Analysis Grid** viewer. You can then horizontally scroll to the **content** column in the **Analysis Grid** to review the operations that were occurring while the selected function was executing, where you might obtain some additional insights into the cause of the debug issue. Lastly, from the **file_line** column of the **Analysis Grid** viewer, you can determine the Samba source code line that initiated logging of the displayed entries, for some further perspective on what may have occurred as the function was executing. Note that you can also drag groups of the **Grouping** viewer into a different position in the group hierarchy to obtain a different analysis perspective on

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		the data. The **SysLogLevels** view **Layout** for the **Chart** viewer enables you to quickly assess the relative distribution of the log entry volumes per **level** value, as derived from your SambaSys log. With this **Layout**, you can obtain an instant visual assessment of the areas in your log that had the most critical levels, which can immediately indicate the direction in which further investigation should proceed. You can also drive selection of log entries in the **Analysis Grid** viewer by double-clicking any bar element of interest in the **SysLog Levels** view **Layout**.
ETW Analysis	.etl	Enable this **Profile** to display the **Analysis Grid** as the default viewer along with the **ETW** view **Layout**, whenever you load data from an event trace log (ETL) file for ETW analysis. You will need to manually open the other viewers that are configured in this **Profile** in the previously described manner, which includes the **Grouping** viewer with the **ETW Guids and IDs** view **Layout**, and **Top Level Protocols Message Count** view **Layout** for the **Chart** viewer. This latter **Layout** uses a **Bar** element visualizer component. **Usage Overview** — the main advantage of the viewer and layout configuration of this **Profile** is that it provides several data sets in different interactive viewing configurations that expose information you will need to quickly isolate problem areas for further investigation of ETW issues. For example, in the **Analysis Grid** viewer you can correlate the **ProcessId** and **ThreadId** that is associated with any event that was logged during execution of a particular process, along with the name of the ETW provider (**Module** column data) that wrote the events that were captured. The **Summary** column data provides additional descriptions, errors, or debug information that can each identify problem areas. **Analysis Example** — as described earlier, applying a `*Summary contains "error"` or `*Summary contains "failure"` Filter from the Filtering Toolbar can

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		be a way to isolate where errors or failures may have occurred. The results of this operation can point you to specific components where errors occurred while also providing a description of what actually occurred. You might also consider executing **Group** commands from the context menus that display when you right-click the headers of the **EventRecord.Header.ProcessId** and **EventRecord.Header.ThreadId** columns, in succession. This will result in a display configuration that organizes the data into groups of events with common **ThreadId** values and nests them under events that have a common **ProcessId** under which the threads executed. The analysis context that this creates can quickly expose which processes carried the highest thread volume, which could be a flag for further investigation. **Note:** The **ThreadId** is a unique identifier of an execution thread that is running under a particular process. The **ProcessId** is a number that is used by the operating system kernel to uniquely identify an active process for which an ETW provider or some other component is generating events. The **Grouping** viewer provides a quick assessment of the event volumes associated each ETW provider that participated in the trace, along with IDs of the events that each provider generated. If you have the ETW manifest for the provider, you may be able to correlate the meaning of events with the IDs that are exposed in any group. You can isolate the events per provider or individual event IDs by clicking a group of interest. If the **Grouping** viewer is in the **Selection Mode** when you click a group, it drives event selection in the **Analysis Grid** viewer. If it is in the **Filtering Mode**, it filters the events into the **Analysis Grid** viewer so that you can analyze additional event **Details**. Note that you can also click the global properties icon on the **Details Tool Window** toolbar for more field information that might be available for a selected event line in the **Analysis Grid** viewer. The fields that are grouped

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
		in the **Grouping** viewer have the following meaning: - **ProviderId** field — specifies the GUID of the ETW trace provider that generated an Event. - **Descriptor.Id** field — specifies the Event identifier, which is part of an Event Descriptor, as described in the ETW Framework Conceptual Tutorial topic. **Tip:** You might also consider selecting the **Process Name and Conversations** layout from the **Layout** drop-down list on the **Grouping** viewer toolbar, to obtain a summary view of all the processes that were initiated across a set of trace results. You can then select a **ProcessName** group of interest to interactively drive the display of corresponding events in the **Analysis Grid** viewer where you can correlate the **ProcessName** with **ProcessId** and **ThreadId** data. This **Profile** contains the **Top Level Protocols Message Count** view **Layout** for the **Chart** viewer. It provides a summary view of the relative distribution of event volumes across a set of trace results for the modules/protocols that generated events in such a trace. This graphic display can immediately point to potential issues where high event volumes are causing large bandwidth consumption.

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
PerfMon Logs	.blg	Enable this **Profile** to display data from a Performance Monitor log and utilize some of Message Analyzer capabilities to manipulate and analyze the data whenever you load data from a *.blg log file. Provides a main display with a graphic representation of performance counter data along with a legend of counters and an adjustable time window for zooming into data points. Displays a related set of messages after you double-click a line of performance counter data for further details. The **Grouping** viewer contains the following groups to organize the data: - Machine - Instance - Counter For any instance, you can click a **Counter** and display that result in the main graphic display. Note that you can double-click a counter data line and display the data that was logged in an associated set of messages in a separate instance of the **Analysis Grid** viewer.

PROFILE NAME	APPLICABLE FILE EXTENSION	DESCRIPTION
NTP Time Offset	.cap	Enable this **Profile** to understand time offset from the network perspective and to troubleshoot time-related issues. The viewer and layout configuration for this **Profile** includes the **NTP Flat** view **Layout** for the **Analysis Grid**; the **NTP Time Offset** view **Layout** for the **Chart** viewer, which shows time offset over time; and the **NTP Source** view **Layout** for the **Grouping** viewer, which organizes the NTP conversations. **Usage Overview** — the main advantage of the viewer and layout configuration of this **Profile** is that you can observe Time Offset data over the timeline of a set of trace results per network conversation, which you can select in a legend to the right of the **Timeline** visualizer component. **Note:** The **Chart** viewer with the **NTP Time Offset** view **Layout** for the **Chart** viewer displays by default for all file types that are associated with this **Profile**. This includes the .cap, .pcap, .etl, and .pcapng file types. Note that the viewer and layout configuration for all these file types is identical in the **Profiles** that apply to them.

More Information

To learn more about the above file types, along with other file types that Message Analyzer supports, see Locating Supported Input Data File Types.

To learn more about the view **Layouts** that you can select for the **Analysis Grid** viewer in a **Profile**, see Applying and Managing Analysis Grid Viewer Layouts.

To learn more about the view **Layouts** that you can select for the **Grouping** viewer in a **Profile**, see Understanding the Built-In Grouping View Layouts.

To learn more about the **Layouts** that you can select for **Chart** viewers in a **Profile**, see Chart Viewer Layouts.

Applying and Managing Profiles

By default, several built-in **Profiles** that exist in the **Advanced Profiles** list are enabled, which means that when Message Analyzer detects that you are loading data from a file type for which a **Profile** has been created and enabled, the **Profile** configuration will be automatically applied after data loading is complete. This action also occurs for any custom-designed and enabled **Profile** of your own. For Message Analyzer to automatically apply any particular **Profile**, the **Use Advanced Profiles** check box must have a check mark in it and the **Profile** must be enabled in the **Advanced Profiles** list. Otherwise, the data viewers and view layouts associated with the **Profile** will not display automatically when you load an associated file type into Message Analyzer. Note that you can enable or disable any **Profile** individually, as described in Enabling and Disabling Profiles.

The remainder of this section describes how to manage **Profiles**, which includes tasks such as enabling or disabling them, creating new **Profiles**, editing **Profiles**, and removing them from the **Advanced Profiles** list.

Enabling and Disabling Profiles

Message Analyzer provides you with the option to either enable or disable any individual **Profile** in the **Advanced Profiles** list on the **Profiles** tab of the **Options** dialog. You can disable a **Profile** by unselecting its check box in the **Enabled** column to the left of the **Profile** name in the **Advanced Profiles** list. This action prevents the **Profile** from activating during the data loading process; however, you can re-enable it at any time by simply placing a check mark back in its check box. You can also disable all **Profiles** simultaneously, even those that are currently *selected*, by removing the check mark from the **Use Advanced Profiles** check box, which prevents Message Analyzer from applying any **Profiles** when you are loading data from a supported file type. To re-enable selected **Profiles**, simply place a check mark back in the **Use Advanced Profiles** check box.

> **NOTE**
>
> If you disable all **Profiles**, Message Analyzer still provides a default **Profile** that specifies the **Analysis Grid** viewer. At your discretion, you can change the default viewer by selecting a new one from the **Default Viewer** drop-down list in the **Default Profile** section on the **Profiles** tab of the **Options** dialog.
>
> This selection determines the default viewer for the display of data in all Live Trace and Data Retrieval Sessions, as described in Session Data Viewer Options. Note that you still have the option of changing the data viewer according to your requirements after you have acquired and displayed session data.

Configuring a New Profile

If you want to create a new **Profile**, you will need to click the **Add Profile** button on the **Advanced Profiles** toolbar on the **Profiles** tab of the **Options** dialog to open the **New Profile** dialog. From here, you can specify the **Profile** configuration that you want by making use of the following controls:

- **Name** — specify a name for the new **Profile**. Be sure to specify a unique name that you can easily recognize and distinguish from other **Profile** names.
- **Description** — optionally specify a short description of the **Profile**.
- **Category** — optionally select a **Category** from this drop-down list. Note that these names are arbitrary and that you can specify a custom category by typing one in the text box portion of this control.
- **File Type** — select one of twenty different supported input file types from the **File Type** drop-down list for your new **Profile**.
- **Copy From** — optionally select one of the **Profiles** in the **Copy From** drop-down list to create an initial pre-populated configuration for your new **Profile** that is based on one of the existing **Profiles**. You will be able to alter the initial configuration of the **Profile** by clicking **Edit Profile** after you **Save** the new **Profile**.
- **Save** — click this button when you are finished with the initial configuration of a new **Profile**.

If you want to make adjustments to the initial configuration that you specified in the **New Profile** dialog, click the **Edit Profile** button on the **Advanced Profiles** toolbar to open a dialog that contains the viewer and layout configuration that you want to modify. From the dialog, you can specify a **Default Viewer** and a view **Layout** for each of the common viewers that all **Profiles** contain, which consist of the following:

- **Analysis Grid** viewer

- **Grouping** viewer
- **Chart** viewer

Example of Configuring a Profile to Create a Targeted Analysis Environment

This section provides an example of creating a **Profile** that specifies data viewers and view layouts that create an environment that uniquely suits analysis of TCP messages. To create this example **Profile**, use the procedure that follows:

1. Display the **New Profile** dialog by clicking the **Add Profile** button on the toolbar above the **Advanced Profiles** list on the **Profiles** tab of the **Options** dialog.

   The **Options** dialog is accessible from the global Message Analyzer **Tools** menu.

2. In the **Name** text box of the **New Profile** dialog, specify a name for your new **Profile** such as "My TCP Analysis".

3. In the **Description** text box of the **New Profile** dialog, optionally specify a brief description of the new **Profile**.

4. In the **Category** drop-down list, optionally specify a category for your **Profile** by selecting one in the list or by typing a custom name in the **Category** combo box.

5. From the **File Type** drop-down list, select the type of file that you want to associate with your new **Profile**, for example, a .cap file.

6. From the **Copy From** drop-down list, select one of the built-in **Profiles** to populate your new **Profile** with initial viewer and view layout settings.

   Note: Use this option if an existing built-in **Profile** contains a configuration from which you want to import settings into your new **Profile**. Otherwise, proceed to the next step.

7. When complete, click the **Save** button in the **New Profile** dialog to save the **Profile**.

8. In the **Advanced Profiles** list, select your newly created **Profile** and then click the **Edit Profile** button on the toolbar above the **Advanced Profiles** list to display the initial configuration of your custom **Profile**.

9. From the **Default Viewer** drop-down list, select the **Analysis Grid** viewer as the default to display your initial session results.

   This list contains the same viewers that are accessible from the **New Viewers** drop-down list on the global Message Analyzer toolbar.

10. From the **File Type** drop-down list, select the **.cap** file type.

    This list contains most of the same file types that are listed in the **All Supported Files** list that displays in the **Open** dialog when you click the **Add Files** button during Data Retrieval Session configuration.

11. From the **Analysis Grid Layout** drop-down list, select the **TCP** view **Layout** for the **Analysis Grid** viewer.

12. From the **Grouping Viewer Layout** drop-down list, select the **TCP Deep Packet Analysis** view **Layout** for the **Grouping** viewer.

13. From the **Charts Layout** drop-down list, select the **TCP Rate and Diagnosis** view **Layout** for the **Chart** viewer.

    As previously described in this topic, this **Chart** will not display unless you select the **Default** item in the **Charts** drop-down list that is accessible from the **New Viewer** drop-down list. You

would typically make this selection after you load data from the .cap file. The **TCP Rate and Diagnosis** view **Layout** for the **Chart** viewer will then display, provided that this **Profile** is enabled in the **Advanced Profiles** list at the time you load the data.

14. Place a check mark in the **Automatically open Grouping Viewer** check box.

    With this check box selected, the **Grouping** viewer will automatically display with populated data in your initial session results when loading data from a .cap file, provided that this **Profile** is enabled in the **Advanced Profiles** list.

15. Click the **Save** button to retain your **Profile** configuration.

At this point, you can create a Data Retrieval Session, as described in Configuring a Data Retrieval Session, to specify a .cap file from which to load data so you can test whether the **Profile** configuration displays the expected default viewer and layouts. The section that follows provides an overview of how you can use the presentation formats of the viewers and layouts of this **Profile** to create some useful analysis contexts.

Targeted TCP Analysis Overview As indicated in the previous procedure, this **Profile** is configured by default to display the **TCP** view **Layout** for the **Analysis Grid** viewer, the **TCP Deep Packet Analysis** view **Layout** for the **Grouping** viewer, and the **TCP Rate and Diagnoses** view **Layout** for the **Chart** viewer. The main advantage of the viewer and layout configuration of this **Profile** is that it provides you with an exceptional context for analysis of TCP messages that can quickly expose potential TCP issues, as described ahead.

TCP Layout — with the **TCP Layout** for the **Analysis Grid** viewer, you can observe values such as **Source** and **Destination** IP addresses, TCP **DestinationPort** and **SourcePort**, **PayloadLength**, **SequenceNumber**, **AcknowledgementNumber**, **WindowScaled**, and a **Summary** description that are each displayed as a separate column of data in the **Analysis Grid** viewer. This provides quick access to important TCP data that can point to areas that need further investigation, for example, an improper receive window size that could be causing packets to be dropped.

TCP Deep Packet Analysis Layout — with the **TCP Deep Packet Analysis** view **Layout** for the **Grouping** viewer, the data displays in a hierarchical grouped configuration that is organized by **DataSource** at the top-level, along with nested groups consisting of the **Network** group for the IP or Ethernet conversations, the **Transport** group that identifies the transport that carried the conversations, and the associated TCP **SourcePort** for each message. Note that you can use the **Grouping** viewer in the **Selection Mode** or **Filtering Mode** to interactively drive the display of messages in the **Analysis Grid** viewer to correlate your data. The **Selection Mode** drives the *selection* of messages in the **Analysis Grid** viewer while the **Filtering Mode** causes *filtered-isolation* of messages in the **Analysis Grid** viewer, where the data displayed in each mode is based on selection of groups in the **Grouping** viewer.

An advantage of the **Grouping** viewer is that it enables you to drill down through the grouped configuration to isolate and expose data of interest at each group level. For example, by clicking a **Network** group, you can interactively select (or filter) all the messages in the **Analysis Grid** viewer that are associated with a particular IP conversation, which is similar to what the Conversation Tree does in the Network Monitor application. In **Selection Mode**, you can analyze the details of selected messages in the context of the original capture sequence, where leading and trailing messages can often provide clues as to why an error might have occurred for a selected message. Another advantage of the **Grouping** viewer is that it can immediately expose the groups that have the highest associated traffic volumes, which can also be a trigger for further investigation.

For further grouping analysis, you might consider using the **Group** command (a right-click command on a chosen **Analysis Grid** viewer column header) to organize the data into separate groups based on common values that exist in the selected column. This feature can quickly expose data that can

enhance your analysis perspective. For instance, you could create a unique analysis context by executing multiple **Group** commands that create a nested group configuration by first grouping the **Module** column and then grouping the **PayloadLength** column. This nested group configuration can quickly expose which **Modules** have messages with the highest payloads so you can isolate such traffic for further investigation. In the grouped context, this could also involve drilling down into the associated message stacks to assess the payload levels, which includes the Transport Layer payloads. With this information and the **Source** and **Destination** address information, you might be able to expose computers that are being overwhelmed by heavy traffic loads where a high volume of TCP retransmits is occurring.

With the **TCP Rate and Diagnoses** view **Layout** for the **Chart** viewer, you can quickly assess how many **Diagnosis** messages occurred in a trace in the context of associated IP conversations, the TCP **SourcePort** and **DestinationPort** associated with the IP conversations, and the ratio (**Rate**) of how many **Diagnosis** messages occurred with respect to the total number of messages in a particular conversation. A high **Diagnosis** error **Rate** can also be a flag that further investigation is warranted. Note that you can obtain a summary of **Diagnosis** message counts and descriptions for each diagnosis type across a set of trace results by opening the **Diagnostics Tool Window** from the global Message Analyzer **Tool** menu.

Taken together, these viewers and layouts provide robust information sets that you can utilize for analysis of TCP data that exists in Network Monitor capture (.cap) files. If you want to modify this **Profile**, you can do so as specified in the section that follows. Because the value of the **ReadOnly** column in the **Advanced Profiles** list for this **Profile** is **False**, you can change the settings of this **Profile** as you wish, going forward.

Editing and Removing Profiles

Given that the **ReadOnly** value for user-created **Profiles** is always set to **False** in the **Advanced Profiles** list, you can edit any **Profile** that you have created at any time, by simply highlighting the **Profile** and then clicking the **Edit Profile** button on the **Advanced Profiles** toolbar. After you modify and save a custom **Profile**, the viewer and layout configuration that you specified will be automatically applied whenever you are loading data into Message Analyzer from the file type for which you configured the **Profile**.

To remove any custom **Profile** that you created, simply highlight the **Profile** in the **Advanced Profiles** list and then click the **Remove Profile** button on the **Advanced Profiles** toolbar. Note that if you delete a custom **Profile**, you will be unable to recover the configuration except by creating a new **Profile**. Note that you cannot **Edit** or **Remove** any of the built-in **Profiles**.

See Also

Analysis Grid Viewer
Grouping Viewer
Chart Viewer Layouts
ETW Framework Conceptual Tutorial

Procedures: Using the Data Viewing Features

43 minutes to read

The procedures in this section demonstrate the use of numerous Message Analyzer features that are described in the Viewing Message Data section. They are intended to serve as a cross section of the many ways in which you can use Message Analyzer viewer features and other integrated functions. As viewing message data and analyzing it are closely related, these procedures demonstrate the use of various viewers and tools that manipulate trace results data for analysis purposes. For additional information on the analysis tools that Message Analyzer provides, see Analyzing Message Data.

> **IMPORTANT**
>
> Although these procedures demonstrate the use of Message Analyzer capabilities in some basic analysis scenarios, they are only a sampling of what you can accomplish with Message Analyzer, given that you can also apply the methodologies described here to many other scenarios.

Procedure Overviews

A brief description of each procedure is included here for review, as follows.

Apply Gradient Style Color Rules — provides an example of how to utilize gradient style **Color Rules** to quickly flag messages that meet the filtering criteria of multiple **Color Rules**.

Apply a Built-In View Layout — provides an example of a built-in view **Layout** that presents a data column configuration that is useful for diagnosing TCP messages when applied, while also automatically grouping messages by IP conversations and ports (**Group** operations on the **Network** and **Transport** columns, respectively) to enhance diagnostic capabilities and perspectives.

Perform Data Grouping Operations — provides several examples of data grouping operations that demonstrate how you can filter and consolidate data from designated **Analysis Grid** viewer columns and reorganize them into separate groups of common properties that greatly enhance your ability to analyze data and resolve issues.

Perform Top-Level Summary Analysis — provides an example of how to use the **Protocol Dashboard** viewer to obtain top-level summaries at a glance for a set of trace results.

Perform Interactive Analysis with Data Viewers — illustrates a simple method for using the **Protocol Dashboard** and **Analysis Grid** viewers together in an interactive manner to enhance data analysis perspectives.

Apply Viewpoints to Trace Data — provides an example that shows you how to use the Message Analyzer **Viewpoints** feature, which enables you to examine data from the viewpoint of a protocol, where the messages of a specific viewpoint protocol are displayed at top-level in the **Analysis Grid** viewer with no message layers above them.

Apply a Time Filter to Trace Results — provides an example that shows you how to apply a **Time Filter** to a set of trace results, so that you can view data in a selected window of time.

Drive Analysis Grid Viewer and Tool Window Interactions — provides an example that demonstrates interaction between the **Analysis Grid** viewer and various tool windows, such as the **Message Data**, **Field Data**, **Message Stack**, **Details**, and **Diagnostics Tool Windows**.

Create an Alias for a Data Field Value — provides an example that demonstrates how to simplify data analysis by creating an **Alias** that substitutes for a cryptic field value.

Create a Union of Two Data Fields — provides an example that demonstrates how to create a **Union** that correlates/combines two data fields with similar values but different names into a single new field that is specified by the **Union** configuration.

Procedures: Using the Data Filtering Features — see this procedural topic for extensive coverage of different ways to apply a view **Filter**.

> **IMPORTANT**
> If you have not logged off Windows after the first installation of Message Analyzer, please log off and then log back on before performing these procedures. This action ensures that in all subsequent logons following installation, your security token will be updated with the required security credentials from the Message Capture Users Group (MCUG). Otherwise, you will be unable to capture network traffic in **Trace Scenarios** that use the **Microsoft-PEF-NDIS-PacketCapture** provider, **Microsoft-Windows-NDIS-PacketCapture** provider, or the **Microsoft-PEF-WFP-MessageProvider**, unless you start Message Analyzer with the right-click **Run as administrator** option.

Apply Gradient Style Color Rules

In the procedure that follows, you will apply the built-in **IPv4 Right Gradient** and **TCP** left gradient **Color Rules** to a Link Layer trace that captured data with the **Local Network Interfaces Trace Scenario** that uses the **Microsoft-PEF-NDIS-PacketCapture** provider (available on Windows 7, Windows 8, and Windows Server 2012 operating systems).

> **NOTE**
> If your machine is running the Windows 8.1 or later operating system, you can capture data in this example with the **Local Network Interfaces Trace Scenario** that uses the **Microsoft-Windows NDIS-PacketCapture** provider.

This procedure demonstrates a simple way to expose TCP messages that have an IPv4 Network Layer in the message stack. The **Color Rule** that is used in the procedure also serves as an example of how you might design multiple gradient-style **Color Rules** with visually-coordinated opposite facing gradients, which you can then use as a troubleshooting mechanism to quickly identify message stack components at a glance.

More Information
To learn more about the concepts upon which this example procedure is based, see Using and Managing Color Rules.

To identify Transport and Network Layer messages with gradient-style Color Rules

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.

   To ensure that you have access to all features, run Message Analyzer as an Administrator.

2. Click **New Session** on the **Start Page** to open the **New Session** dialog.

3. Under **Add Data Source** in the **New Session** dialog, click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features that it contains in the **New Session** dialog.

4. In the **Network** category of the **Select Scenario** drop-down list on the **Live Trace** tab, click the **Local Network Interfaces Trace Scenario**.

   If your operating system is Windows 7, Windows 8, or Windows Server 2012, the **ETW Providers** list on the **Live Trace** tab is populated with the **Microsoft-PEF-NDIS-PacketCapture** provider **Name** and **Id** (GUID). Otherwise, for the Windows 8.1, Windows Server 2012 R2, Windows 10 operating system, or later, the **Microsoft-Windows-NDIS-PacketCapture** provider information displays.

5. Click the **Start** button in the **New Session** dialog to automatically select the default data viewer and start capturing data. Assuming that you have not changed the default data viewer in the **Default Profile** pane on the **Profiles** tab of the global **Options** dialog, the default viewer will be the **Analysis Grid**.

   Message Analyzer should immediately begin capturing data and accumulating it in the **Analysis Grid** viewer.

6. While Message Analyzer is capturing data, attempt to reproduce any conditions that are related to a particular TCP or IPv4 issue you are trying to isolate, for example, network connection or packet loss problems.

7. Stop the trace at a suitable point by clicking the **Stop** button on the global Message Analyzer toolbar.

> **TIP**
> You can temporarily suspend tracing operations by clicking the **Pause/Resume** button and you can resume tracing by clicking the **Pause/Resume** button again.

8. Click the **Color Rules** drop-down list on the **Analysis Grid** viewer toolbar, and then under the **Network** category of the drop-down that displays, select the **TCP** left gradient and **IPv4 Gradient Right Color Rules**.

   All top-level TCP messages or other top-level messages that have TCP in the origins tree are highlighted with the light blue left-to-right gradient **Color Rule** style. Also, all TCP messages that have an IPv4 network layer are highlighted in the olive green right-to-left gradient **Color Rule** style, thus enabling you to easily view messages that meet the filtering criteria of both the applied **Color Rules**.

> **NOTE**
> To isolate either TCP or IPv4 messages at top-level to further enhance analysis, you can apply a TCP or IPv4 viewpoint as appropriate from the Filtering toolbar.

Apply a Built-In View Layout

In the procedure that follows, you will apply the built-in **TCP Deep Packet Analysis with Absolute Sequence Number Grouping** view **Layout** to trace data that is displayed in the **Analysis Grid** viewer. This view **Layout** has a column layout configuration that contains various TCP fields, the values of which can be important when diagnosing TCP issues. The columns that hold TCP field data include **DestinationPort**, **SourcePort**, **PayloadLength**, **SequenceNumber**, **AcknowledgementNumber**, and **WindowScaled** columns. In addition, a **TimeDelta** field is also included to display the running time for captured messages. The predefined **Layout** also includes groupings of **Network** and **Transport** columns that present the details of the IP conversations that took place on corresponding TCP ports within a trace. Note that the **Network** and **Transport** columns were removed after the **Grouping** operation, but before this **Layout** was saved in the default **Layout** Library item collection.

More Information
To learn more about the concepts upon which this example procedure is based, see Applying and Managing Analysis Grid Viewer Layouts.

To apply a built-in View Layout for TCP diagnosis

1. Perform steps 1 through 7 of the procedure To identify Transport and Network Layer messages with gradient-style Color Rules to start and stop a Message Analyzer Live Trace Session that uses the **Local Network Interfaces Trace Scenario**.

2. Click the **Layout** drop-down list on the **Analysis Grid** toolbar and then click the **TCP Deep Packet**

Analysis with Absolute Sequence Number Grouping item in the drop-down list that appears.

The new column configuration displays and the data is grouped into **Network**, **Transport**, and **Sourceport** groups, as indicated by corresponding labels above the tree grid. The data groups are also organized such that the **Transport** nodes are nested within the **Network** nodes, and the **Sourceport** nodes are nested within those. Note that the **Network** conversations can use either IP or Ethernet addresses.

3. Expand a particular **Network** node to expose the **Transport** node it contains.

   The exposed **Transport** node provides an indication of the number of messages that it contains, along with the source and destination TCP ports over which IP or Ethernet conversations took place.

4. Expand the **Transport** node to display the TCP messages, so that you can examine the TCP field data. If you are dealing with loss of packets, you might check the **WindowScale** field for low values.

   > **NOTE**
   > For convenience of viewing the TCP column data, you can alter the data columns that will horizontally scroll by right-clicking the column that you want as the first scrollable column and then select the **Freeze Columns to Left** command in the context menu that appears.

5. Expand the **Sourceport** nodes, so that you can view the messages that transited each TCP port.

6. Repeat steps 3, 4, 5 for other **Network**, **Transport**, and **Sourceport** nodes as appropriate.

7. To obtain a different perspective on the data, you can drag the **Transport** group label and drop it to the left of the **Network** group label.

   The data is now grouped and organized with **Network** nodes nested within the **Transport** nodes.

   Note: You can drag and drop any of the Group labels into any position that you want, to change the way messages are hierarchically organized.

> **NOTE**
> To restore the default **Layout**, click the **Layout** drop-down list on the **Analysis Grid** viewer toolbar, click the **Manage Layouts** item, and then click the **Restore Application Default Layout** command that displays in the submenu that appears.

Perform Data Grouping Operations

In the procedure that follows, you will execute the **Group** command on various **Analysis Grid** viewer data columns, including the **ContentType**, **Transport**, **Source** or **Destination**, and **Diagnosis** data columns. The grouping operations will enable you to quickly determine the object types being requested by your web browser, assess the heaviest port traffic, determine the IP addresses carrying the most traffic, and examine grouped diagnosis messages types, respectively. In this procedure, the **Analysis Grid** viewer will be populated with message data that you capture with the **Microsoft-PEF-WFP-MessageProvider** in the **Loopback and Unencrypted IPSEC Trace Scenario** and the focus will be on Application Layer (HTTP) and Transport Layer messages.

More Information
To learn more about the concepts upon which this example procedure is based, see Using the Analysis Grid Group Feature.

To perform multiple data grouping operations for analysis

1. Perform steps 1 through 3 of the procedure To identify Transport and Network Layer messages with

gradient-style Color Rules, to start Message Analyzer and open the **New Session** dialog for Live Trace Session configuration.

2. In the **Network** category of the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, click the **Loopback and Unencrypted IPSEC Trace Scenario**. Alternatively, click the **Loopback and Unencrypted IPSEC Trace Scenario** in the **Favorite Scenarios** list that is accessible from the Message Analyzer **File** menu.

   The **Microsoft-PEF-WFP-MessageProvider** information displays in the **ETW Providers** list on the **Live Trace** tab, which includes the provider **Name**, **GUID**, and a **Configure** link that opens the **Advanced Settings** dialog for this provider.

   > **NOTE**
   > In addition to capturing loopback traffic and unencrypted IPSEC messages, the **Microsoft-PEF-WFP-MessageProvider** minimizes other lower-level noise such as broadcast traffic, so that you can focus your analysis above the IP/Network Layer. Also note that messages below the Transport Layer are typically represented in the message stack as a **WFPCapture** and below that are events at the **ETW** level.

3. If the **Start With** drop-down list in the **New Session** dialog does not indicate the **Analysis Grid** viewer, then click the drop-down list and select it.

4. Click the **Start** button in the **New Session** dialog to start capturing data.

   Message Analyzer may immediately begin capturing and accumulating message data in the **Analysis Grid** viewer.

5. While Message Analyzer is capturing data, launch a web browser and attempt to reproduce any conditions that are related to a particular HTTP issue you might be experiencing. For example, you might attempt to navigate to a poorly performing web server with your browser.

6. Stop the trace at a suitable point by clicking the **Stop** button on the global Message Analyzer toolbar.

7. Click the **Add Columns** button on the **Analysis Grid** viewer toolbar to display the **Field Chooser Tool Window** in focus in its default location, if it is not already displayed.

8. Open the **HTTP** node in **Field Chooser** and navigate to the **ContentType** field in the HTTP **Operation** message hierarchy, right-click the field, and then select the **Add As Column** context menu item to add the **ContentType** column to the **Analysis Grid** viewer.

9. Open the **TCP** node in **Field Chooser** and navigate to the **Transport** field in the **Segment** message hierarchy, right-click the field, and then select the **Add As Column** context menu item to add the **Transport** column to the **Analysis Grid** viewer.

10. In a similar manner, add the **ResponseTime** field from the **Global Annotations** node in **Field Chooser** as a new column in the **Analysis Grid** viewer, so that you can determine how quickly the web server is responding to HTTP requests.

11. Right-click the **ContentType** column in the **Analysis Grid** viewer and select the **Group** command from the context menu.

    The trace data is grouped according to the different content types associated with HTTP messages, so that you can examine the types of objects being passed to your web browser by the server.

TIP

To create a more focused analysis, you can limit the display to HTTP messages only by specifying an HTTP **Viewpoint**; you can do this by clicking the **Viewpoint** drop-down list on the Filtering toolbar and then selecting the **HTTP** item.

When you are done with assessing the data, click the **x** in the **ContentType** Group label above the **Analysis Grid** to remove the Group and return to the default view **Layout**.

12. Sort the **ResponseTime** column in descending order and then **Group** this column to quickly expose the slowest server responses and associated messages in separate Groups for analysis.
13. Remove the **ResponseTime** Group by clicking the **x** in the **ResponseTime** Group label above the **Analysis Grid**.
14. Right-click the **Transport** column in the **Analysis Grid** viewer and select the **Group** command from the context menu.

    The trace data is grouped according to the different Transport types, such as TCP or UDP, so that you can examine the ports across which the most substantial traffic is transiting.
15. Remove the **Transport** grouping in the previously indicated manner and then execute the **Group** command on the default **Source** or **Destination** column of the **Analysis Grid** viewer.

    The trace data is grouped according to **Source** or **Destination**, as appropriate, so that you can determine which IP addresses are carrying the most traffic. Note that you can obtain similar statistics by executing a **Group** command on the **Network** column, which you can add from the IP message hierarchy in the **Field Chooser**.

TIP

You can also nest groups by performing successive **Group** operations on multiple columns. For example, you might also add the **Response.PayloadLength** field from the **HTTP** message hierarchy in **Field Chooser** and then perform multiple groupings in succession on the **ResponseTime**, **ContentType**, **PayloadLength**, and **Transport** data columns so you can view the slowest server response times correlated with the content types that are associated with the highest payload values, along with the TCP ports that carried that information. Another grouped configuration you can try would be to **Group** the **Source** and **Destination** columns, in that order, to organize all the **Destination** traffic that is associated with each **Source** address, or vice versa.

16. Remove all **Group** configurations to return to the original **Analysis Grid** viewer display and then execute the **Group** command on the **DiagnosisTypes** column in the **Analysis Grid** viewer.

    The trace data is grouped according to the different types of diagnosis messages, which includes **Application**, **InsufficientData**, **Parsing**, and **Validation** message types, so that you can immediately assess the types of errors that occurred in your trace. For more information about the meaning of these diagnosis message types, see the "Enum Values for DiagnosisType filters" table in the Diagnosis Category topic.

Perform Top-Level Summary Analysis

In the procedure that follows, you will use several viewing infrastructure components to accomplish simple data analysis tasks. For example, you will use the graphic chart visualizer components of the **Protocol Dashboard** viewer, which includes the **Top Level Protocol Summary Bar** element, **Pie** chart, **Table** grid, and the **Timeline** (**Top Level Protocols Over Time**) visualizer components, to view top-level protocol summary data that can reflect traffic volume levels for the message types in a trace, along with message activity across selected windows of time into which you can zoom. In the first part of the procedure, you will use the **Microsoft-PEF-NDIS-**

PacketCapture provider in the **Local Network Interfaces Trace Scenario** to capture message data in a Live Trace Session. However, if you are running the Windows 8.1 or Windows Server 2012 R2 operating system, you will be using the **Microsoft-Windows-NDIS-PacketCapture** provider in this **Trace Scenario**.

This example also shows how to use the **SMB Reads and Writes Bytes/Second**, **SMB File Stats**, and **SMB/SMB2 Service Performance** view **Layouts** for **Charts** to expose file access activities and statistics. In this part of the procedure, you will start a new session with the **Loopback and Unencrypted IPSEC Trace Scenario**, in which you will use the **Microsoft-PEF-WFP-MessageProvider** to focus on statistical summaries of SMB/SMB2 file access operations messages above the IP/Network Layer.

More Information
To learn more about the **Protocol Dashboard** viewer, see the Protocol Dashboard topic.
To learn more about the SMB **Layouts** for the **Chart** viewer, see the subtopics in the File Sharing Category section.

To analyze top level summary data

1. Perform steps 1 through 5 of the procedure To identify Transport and Network Layer messages with gradient-style Color Rules to start a Message Analyzer Live Trace Session with the **Local Network Interfaces Trace Scenario**.

2. While Message Analyzer is capturing data, attempt to reproduce conditions that are related to any particular issue you might be trying to resolve, for example, a high volume of TCP traffic to a target computer.

3. Stop the trace at a suitable point by clicking the **Stop** button on the global Message Analyzer toolbar.

4. From the **New Viewer** drop-down list on the global Message Analyzer toolbar, click **Charts (Deprecated)** and then select the **Protocol Dashboard** item in the drop-down list.

5. In the **Protocol Dashboard** viewer, observe the numerical and graphical presentation of top protocol activity in the trace by examining the **Top Level Protocol Summary** and **Top Level Protocols Over Time** visualizer components, which expose the relevant statistics.

   In the **Top Level Protocol Summary** Table and Bar element sections of the dashboard, you can observe message traffic volume that is sorted in a descending scale from highest to lowest. Note that an extraordinarily high traffic volume for a particular module can immediately expose the top bandwidth consumer, or exceptionally heavy TCP traffic might indicate that you have a large quantity of TCP retransmits or duplicate ACK messages in your trace.

   If you suspect there is an issue with a protocol or module that has particularly high traffic volume, you can double-click the Bar element or Pie chart segment representing the module in the **Top Level Protocol Summary** to display only those specific messages in a separate instance of the **Analysis Grid** viewer for further investigation. You can also adjust the time window slider controls of the Timeline visualizer component (**Top Level Protocols Over Time**) to zoom into specific messages in a particular time slot and then double-click a message node to display that traffic only in a separate instance of the **Analysis Grid** viewer for further investigation.

6. Start another Live Trace Session with the **Loopback and Unencrypted IPSEC Trace Scenario** and capture data live with Message Analyzer while performing file access operations that have previously been problematic.

7. Stop the trace at a suitable point by clicking the **Stop** button on the global Message Analyzer toolbar and then launch the **SMB Reads and Writes Bytes/Second** view **Layout** from the **Chart** drop-down in the **New Viewer** drop-down list on the global Message Analyzer toolbar.

   From this view **Layout**, you can obtain statistics that reflect the network bandwidth being consumed by the file access/sharing activities of the Server Message Block (SMB) protocols. See SMB Reads and Writes

Bytes/Second for further details.

8. Adjust the time window slider controls in the **SMB Reads and Writes Bytes/Second** viewer to zoom into specific messages in a particular time slot.

9. Double-click a message node or time line in the **SMB Reads and Writes Bytes/Second** data viewer to display specific traffic in a new instance of the **Analysis Grid** viewer tab for further investigation.

> **TIP**
> You can also use the **Field Chooser Tool Window** to add an SMB **FileName** or SMB2 **FileName** column to the **Analysis Grid** viewer and then execute a **Group** command on the new column so that you can examine the SMB traffic that is associated with access to specific files.

10. Optionally, select the **SMB File Stats** view **Layout** from the **Chart** drop-down in the **New Viewer** drop-down list. With this **Layout**, you can examine a summary of SMB file statistics in a **Table** visualizer component that includes the file name, access duration, total number of bytes for each file or folder access operation, and the data transmission rates, as described in SMB File Stats.

    You might also consider selecting the **SMB/SMB2 Service Performance** viewer **Layout**, also in the **Chart** drop-down, to examine statistics that expose how long first responses to SMB operations are taking (**ResponseTime**), possibly to expose slow server response issues; and how long it is taking for operations to complete (**ElapsedTime**), as a possible indication of network issues; as described in SMB/SMB2 Service Performance.

Perform Interactive Analysis with Data Viewers

The procedure that follows provides a simple example of how you might utilize the **Analysis Grid** and **Protocol Dashboard** viewers interactively to analyze captured message data. The example also indicates how you might user other data viewers interactively with the **Analysis Grid** viewer:

More Information
To learn more about data viewers and how they interact, see Data Viewer Concepts.

To analyze data through data viewer interaction

1. Perform steps 1 through 4 of the procedure To identify Transport and Network Layer messages with gradient-style Color Rules, to start Message Analyzer, open the **New Session** dialog for Live Trace Session configuration, and select the **Local Network Interfaces Trace Scenario**.

2. Click the **Start** button in the **New Session** dialog to begin capturing data in a Live Trace Session.

    The captured data begins to accumulate in the **Analysis Grid** viewer, assuming that you have not changed the default data viewer in the **Default Profile** pane on the **Profiles** tab of the global **Options** dialog that is accessible from the global Message Analyzer **Tools** menu.

3. While Message Analyzer is capturing data, attempt to reproduce conditions that are related to any particular issue you are trying to isolate.

4. Stop the trace at a suitable point by clicking the **Stop** button on the global Message Analyzer toolbar.

5. From the **New Viewer** drop-down list on the global Message Analyzer toolbar, click **Charts (Deprecated)** and then select the **Protocol Dashboard** item in the drop-down list.

6. In the **Protocol Dashboard** viewer, observe the relative distribution of captured message volumes in the **Top Level Protocol Summary** Bar element visualizer of the **Protocol Dashboard**, in an attempt to isolate message traffic that might be related to failures in a particular component, system, or service. Note

that high message volumes can be a flag for underlying issues that may need further investigation, such as network issues and TCP retransmits.

7. In the Bar element visualizer, double-click the graphic bar representing the message traffic you want to target, for example, a protocol that has a high volume of messages.

   A new instance of the **Analysis Grid** viewer opens and contains only the traffic that you targeted, for further analysis.

8. Sort the **DiagnosisTypes** column in the **Analysis Grid** viewer to bubble up any errors that might have occurred in the target traffic.

9. Perform a **Group** operation by right-clicking the **DiagnosisTypes** column in the **Analysis Grid** viewer and then selecting the **Group** command from the context menu that displays. The data is then organized into expandable group nodes that each contain different diagnosis message types. By expanding each node, you can view the messages that contain the diagnosis errors.

10. Click the diagnosis error icons in the **DiagnosisTypes** column under the expanded group nodes to review the error message text. You might also examine the **Summary** column descriptions for these messages in the **Analysis Grid** viewer to discover any evidence of the underlying failures that are associated with the diagnosis errors that occurred. You can also review the values of the fields for selected messages in the **Details Tool Window**.

More Information
To learn more about the meaning of diagnosis message types, see the "Enum Values for DiagnosisType filters" table in the Diagnosis Category topic.

Apply Viewpoints to Trace Data

In the procedure that follows, you will apply HTTP and TCP **Viewpoints** so that you can view HTTP- or TCP-related traffic at top-level, without having to drill down into each message stack to expose these messages. In addition, you will alternately disable or enable Operations so that you can expose HTTP request and response messages — either in their original chronological order (by executing the **Disable Operations Viewpoint**), or encapsulated by Message Analyzer in top-level Operation rows (by executing the **No Viewpoint** command) respectively, in the **Analysis Grid** viewer.

More Information
To learn more about the concepts upon which this example procedure is based, see Applying and Managing Viewpoints.

To analyze data with applied Viewpoints

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.

   To ensure that you have access to all features, run Message Analyzer as Administrator.

2. Click the Message Analyzer global **File** menu and then click **New Session** to display the **New Session** dialog.

3. Under **Add Data Source** in the **New Session** dialog, click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features that it contains in the **New Session** dialog.

4. In the **Network** category of the **Select Scenario** drop-down list on the **Live Trace** tab, click the **Loopback and Unencrypted IPSEC Trace Scenario**.

   The **Microsoft-PEF-WFP-MessageProvider** information displays in the **ETW Providers** list on the **Live**

Trace tab, which includes the provider **Name**, **GUID**, and a **Configure** link that opens the **Advanced Settings** dialog for this provider.

> **NOTE**
> In addition to capturing loopback traffic and unencrypted IPSEC messages, the **Microsoft-PEF-WFP-MessageProvider** minimizes other lower-level noise such as broadcast traffic, so that you can focus your analysis above the IP/Network Layer.

5. If the **Start With** drop-down list in the **New Session** dialog does not indicate the **Analysis Grid** viewer, then click the drop-down list and select it.
6. Click the **Start** button in the **New Session** dialog to start the trace and begin capturing data.

   The captured data begins to accumulate in the **Analysis Grid** viewer.

> **NOTE**
> You can start this type of trace immediately with default provider configuration settings, by selecting the **Loopback and Unencrypted IPSEC Trace Scenario** in the **Favorite Scenarios** list on the Message Analyzer **Start Page**. By starting a trace in this manner, you will not have access to provider configuration settings prior to capturing data.

7. While Message Analyzer is capturing data, launch a web browser and attempt to reproduce any conditions that are related to a particular HTTP issue you are trying to isolate, for example, a slowly responding or non-responsive web server.
8. Stop the trace at a suitable point by clicking the **Stop** button on the global Message Analyzer toolbar.
9. In the **Analysis Grid** viewer, note that you have HTTP messages displaying as top-level operation message rows, as signified by messages with a blue-cubed icon to the left of the message number, along with some TCP messages at top-level and others hidden within the message stack. You may also have HTTP fragments hidden within **Analysis Grid** viewer expansion nodes. This configuration of displayed messages is typical of the results returned by the **Microsoft-PEF-WFP-MessageProvider**, which focuses on Transport Layer messages and above.
10. Click the **Viewpoints** drop-down list on the Filtering toolbar and select the **HTTP** item from the list.

    All HTTP messages are driven to top-level message rows in the **Analysis Grid** viewer, which can also include any fragments that existed in the origins tree (message stack). In this view configuration, you can focus on HTTP messages without the encumbrance of other message types in display. However, because associated HTTP request and response messages are still grouped as Operations to provide context, there is a chronological displacement of response messages in this configuration that you can resolve by disabling the Operations. See the **Important** note below.
11. Click the **Disable Operations** item in the **Viewpoints** drop-down list on the Filtering toolbar just below the **Analysis Grid** viewer tab.

    Note that the HTTP request and response message pairs that were formerly grouped under Operation nodes are now displayed in chronological order in the **Analysis Grid** viewer. This view should be familiar to Network Monitor users and should enable them to work with the messages in the manner with which they are accustomed. However, for quick analysis of HTTP request and response pairs, it is more expedient to view the data encapsulated in the default Operation node format to see the information at top-level. At this level, you can still use typical data analysis tools such as sorting, grouping, filtering, and so on.

    Also note that data values that are important to HTTP analysis include **ResponseTime** and **ElapsedTime**, as described in the **Important** note below.

12. To return to the default display, click the **No Viewpoint** item in the **Viewpoints** drop-down list on the Filtering toolbar to return to the default viewpoint. This is the presentation format in which the **Analysis Grid** viewer normally displays.

13. Next, isolate TCP messages at top-level in the trace by selecting the **TCP** item in the **Viewpoint** drop-down list.

    All TCP messages display in top-level message rows in the **Analysis Grid** viewer. Note that you will not see any Operation message nodes in this view because the **TCP Viewpoint** filters out everything above it.

14. From the **TCP Viewpoint**, use your typical data analysis tools such as sorting, grouping, filtering, pattern matching, and annotating, in addition to viewing message **Details**, **Message Stack**, and **Diagnostics Tool Window** data, to analyze the information. You might also apply the built-in **TCP Deep Packet Analysis with Absolute Sequence Number and Grouping Layout** from the **Layout** drop-down list on the **Analysis Grid** toolbar, so you can focus on important TCP field data.

> **IMPORTANT**
> Pairing up request and response messages in Operation nodes for protocols that typically use request/response pairs such as HTTP, DNS, and SMB2, provides immediate access to response messages rather than having to search through potentially hundreds or even thousands of messages to find them. Another advantage of this configuration is that you can readily measure and correlate important values such as **ResponseTime** and **ElapsedTime**, which specify how long it took for the first server response and how long it took to receive all message fragments and complete the Operation, respectively. High values for these times can provide an indication of a poorly responding server in the first case and network latency issues in the second. The **ElapsedTime** is displayed by default in the **Analysis Grid** viewer column layout; however, you must add the **ResponseTime** column by right-clicking it under **Global Annotations** in the **Field Chooser Tool Window** and then selecting the **Add as Column** command.

> **More Information**
> **To learn more** about average response time for Operations, see Average Response Time for Operations.
> **To learn more** about average elapsed time for Operations, see Average Elapsed Time for Operations.

Apply a Time Filter to Trace Results

In the procedure that follows, you will start a Data Retrieval Session and load data from a saved trace or log file. You will then apply a **Time Filter** to the loaded message collection so that you can temporarily focus on analyzing messages in a specified window of time. You will also verify that you can toggle back and forth between the time-filtered data and your original data, as your analysis might require.

More Information
To learn more about the concepts upon which this example procedure is based, including the benefits of using a **Time Filter**, see Applying a Time Filter to Session Results.

To apply a Time Filter to a set of trace results

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.

   To ensure that you have access to all features, run Message Analyzer as an Administrator.

2. Click the global Message Analyzer **File** menu and then click **New Session** to display the **New Session** dialog.

3. Under **Add Data Source** in the **New Session** dialog, click the **Files** button to display the **Files** tab along with the associated session configuration features that it contains in the **New Session** dialog for a Data Retrieval Session.

4. On the **Files** tab of the **New Session** dialog, click the **Add Files** button on the **Files** tab toolbar to launch the **Open** dialog, select a large trace or log file containing data that you want to view in a specific time window, and then click **Open**.

The name of the trace or log file appears in the files list.

5. If you loaded a *.log file, you should choose an applicable configuration file from the **Text Log Configuration** drop-down list just below the toolbar of the **Files** tab for your log, to enable full parsing of messages.

> **IMPORTANT**
> A built-in OPN configuration file is required to parse a text log. Message Analyzer provides several configuration files by default, including Cluster, IIS, SambaSys, and Netlogon .log files, among many others. If the type of text log you want to parse is not included in the previously specified list, you will need to create your own custom OPN configuration file to parse your log, as described in Addendum 1: Configuration Requirements for Parsing Custom Text Logs.

6. Confirm that the **Analysis Grid** viewer is selected in the **Start With** drop-down list and then click the **Start** button in the **New Session** dialog to start loading the message data.

> **IMPORTANT**
> Do *not* configure a **Time Filter** in the Data Retrieval Session configuration, given that you will be doing this only *after* the data is loaded into Message Analyzer. You might do this in a different scenario where you want to improve performance by limiting the amount of input messages from a high volume data file, as described in Considering Performance vs. Usability Factors for Time Filter Application, but not in this particular example.

7. After the message data is loaded and displayed in the **Analysis Grid** viewer, click the **Add Time Filter** item from the **Add Filter** drop-down list on the Filtering toolbar to open the **Time Filter** panel.
8. In the **Time Filter** panel, use the **Start Time** and **End Time** slider controls to configure a window of time in which you want to view data. As you do this, you will see the **Start Time** and **End Time** values change.
9. When you are done with **Time Filter** configuration, click the **Apply** button in the **Time Filter** panel to filter the message data according to the time window that you specified.

   The number of messages displaying in the default data viewer is reduced in accordance with the specified **Time Filter** configuration, thus enabling you to perform analysis on a focused data set.

   Note that the number of messages that display are indicated next to the **Available** label on the Message Analyzer status bar at the bottom of the UI.
10. To remove the time filtering configuration that you applied, click the **Remove** button in the **Time Filter** panel to return to your original data.
11. To reapply the same time filtering configuration, click the **Apply** button again.

    Note that you can toggle the application and removal of a **Time Filter** as many times as your analysis requires. You can also apply view **Filters**, **Viewpoints**, and **Viewpoint Filters** to the time filtered results to further focus the results you want to analyze.

> **NOTE**
>
> **Time Filters** do not persist across sessions or even across viewers of the same session, which means that you will need to create a new **Time Filter** configuration for every session or viewer where you want to apply time window filtering. Also note that you can save any data set to which you have applied a **Time Filter** by clicking **Save As** on the global Message Analyzer **File** menu and specifying the **Filtered Messages for Analysis Grid view** option to perform the save with the **Save/Export Session** dialog.

Drive Analysis Grid Viewer and Tool Window Interactions

In the procedure that follows, you will run a trace and display data in the **Analysis Grid** viewer. Thereafter, through message selection in the **Analysis Grid** viewer, or message and field selection in various **Tool Windows**, the procedure will demonstrate how to interactively drive the display of data in these viewing components to facilitate rapid assessment of message details, which include field values and types, hexadecimal or binary data, diagnosis message types and details, message stack configurations, and so on. This procedure assumes that certain tool windows you will be working with are not currently displayed in the Message Analyzer analysis surface. If they are already displayed, please ignore the steps that specifically require you to display them.

More Information

To learn more about how to position Message Analyzer data viewers and **Tool Windows** for enhanced data analysis, see Working with Message Analyzer Window Layouts.

To drive interaction between the Analysis Grid viewer and tool windows

1. Perform steps 1 through 7 of the procedure To identify Transport and Network Layer messages with gradient-style Color Rules to start and stop a new Message Analyzer Live Trace Session that uses the **Local Network Interfaces Trace Scenario**.
2. Click the global Message Analyzer **Tools** menu, click the **Windows** item, and then select **Diagnostics** in the submenu that appears, to display the **Diagnostics Tool Window** in its default docking location.

> **NOTE**
>
> The **Diagnostics** window is a preview feature that will not be included in the **Windows** submenu unless you have first selected it on the **Features** tab of the **Options** dialog. This dialog is also accessible from the global Message Analyzer **Tools** menu. Note that a Message Analyzer restart is required after this selection.

3. Click the global **Tools** menu again, click the **Windows** item, click the **Message Stack** item, and then select **Message Stack 1** in the submenu to display the **Message Stack 1 Tool Window** in its default docking location.
4. Ensure that the **Message Stack 1** window is in focus (click its tab) and then select any message in the **Analysis Grid** viewer.

   Observe that message selection in the **Analysis Grid** viewer drives message selection in the **Message Stack 1** window and message details in the **Details Tool Window**.
5. Click the **Diagnostics** window tab in its default docking location to bring it into focus and then select one or more diagnosis message types in the **Diagnostics** grid.

   Observe that message selection in the **Diagnostics** window drives selection of one or more top-level messages in the **Analysis Grid** viewer and corresponding message details in the **Details** window.

> **NOTE**
>
> You should be aware that even though top-level messages are highlighted, the actual message that contains a diagnosis error might be at a lower layer in the origins tree (message stack). You can determine this by expanding message nodes in the **Analysis Grid** viewer under the highlighted top-level message/s. In addition, note that the **Diagnostics** window data columns provide descriptions of all diagnosis messages in the current trace results, so you do not have to drill down into the origins tree through node expansion to see them.

6. Click the **Message Data 1 Tool Window** tab in its default docking location to bring it into focus and then select any message in the **Analysis Grid** viewer.

   Observe that message selection in the **Analysis Grid** viewer drives the display of hexadecimal, binary, or ASCII data selection in the **Message Data 1** window and message details in the **Details** window.

> **NOTE**
>
> If you select the top-level message in any Operation row, it does not display any data in any **Message Data** window. Rather, you must expand the Operation node in the **Analysis Grid** viewer and select one of the nested request or response messages that it contains to display highlighted hexadecimal data.

7. Select any message in the **Analysis Grid** viewer and then select a field **Name** in the **Details** window.

   Observe that field selection in the **Details** window drives the display of a hexadecimal, binary, or ASCII field value in any **Message Data** window and a field value in the **Field Data Tool Window** as well, provided that the field you selected in **Details** had a field *value*.

Create an Alias for a Data Field Value

In the procedure that follows, you will perform a live trace and display the results data in the **Analysis Grid** viewer. You will then create an **Alias** for an IPv6 address and name it with a string value of "MyComputer". You will then use the new **Alias** in a Filter Expression that you apply to the trace.

More Information

To learn more about the concepts upon which this example procedure is based, see Using and Managing Message Analyzer Aliases.

To create a field value Alias

1. Perform steps 1 through 7 of the procedure To identify Transport and Network Layer messages with gradient-style Color Rules to start and stop a new Message Analyzer Live Trace Session that uses the **Local Network Interfaces Trace Scenario**.

2. In the **Destination** column of the **Analysis Grid** viewer, right-click an IPv6 address for the local computer and select the **Create Alias for 'Destination'...** item in the context menu that displays. The **Alias Editor** dialog displays, in which you can specify an **Alias** name, **Description**, and **Category**.

   If you do not know the IPv6 address of the local computer, run `IPConfig /All` at the command line to determine it.

3. In the **Alias** text box of the **Alias Editor** dialog, specify a friendly name such as "MyComputer", or specify another name that is appropriate for your environment.

4. In the **Description** text box of the **Alias Editor** dialog, enter a description that identifies the purpose of the **Alias**, for future reference and for identification if you intend to share the **Alias** with other users.

   The **Description** text will display in a tool tip when you hover over the **Alias** name in the **Aliases** drop-

down list with your mouse, or when you hover over the **Alias** name in the **Manage Alias** dialog.

5. In the **Category** combo box of the **Alias Editor** dialog, either select an existing **Category** or specify a new one, for example "IPv6 Addresses".

   Any new **Category** that you specify appears as a subcategory under the top-level **My Items** category and will contain the new **Alias** after you **Save** it.

6. In the **Alias Editor** dialog, ensure that the **Auto Refresh Views** check box is selected if you want Message Analyzer to immediately perform a refresh of all data viewers that will be impacted by application of the new **Alias**.

7. In the **Alias Editor** dialog, click the **Save** button to save your new **Alias**.

   All data viewers, including the **Analysis Grid** viewer and any **Chart** viewer **Layout** that is displayed, are updated to reflect application of the new **Alias**, providing that the **Auto Refresh Views** check box was selected when you saved the **Alias**. If this is the case, observe that the IPv6 address of the local computer is now identified in the **Source** and **Destination** address columns of the **Analysis Grid** viewer as "MyComputer".

   Also verify that the new **Alias** appears in the **Aliases** drop-down list, which is accessible from the global Message Analyzer **Tools** menu. In a similar manner, the **Alias** should also appear in the **Manage Alias** dialog.

8. In the text box on the default **Filter** panel of the Filtering toolbar, enter the following text to create a Filter Expression that uses your new **Alias**:

   ```
   *Source == "MyComputer"
   ```

> **NOTE**
> If a **Filter** panel on the Filtering toolbar is not currently displayed, click the **Add Filter** drop-down list just below the **Analysis Grid** viewer tab and select the **Add Filter** item in the list.

9. Click the **Apply** button on the **Filter** panel of the Filtering toolbar and observe that the specified filter removes all traffic except the messages in which "MyComputer" represents either the **Source** or **Destination** computer IPv6 address.

> **NOTE**
> If you want to save the filter you created in this procedure, select the **New Filter** item in the **Library** drop-down on the **Filter** panel of the Filtering toolbar and provide a **Name**, **Description**, and **Category** for the filter in the **Edit Filter** dialog that appears before you **Save** it. Note that the **Edit Filter** dialog automatically captures the Filter Expression text that you specified in the **Filter** panel text box.

Create a Union of Two Data Fields

In the procedure that follows, you will load data into Message Analyzer from two saved files that contain related data that was captured in a common environment and within the same time frame; one from a log file and the other from a former live trace. The procedure specifies files that contain messages from SMB operations that have identical value data for certain fields, but which are named differently. After you load your data files into Message Analyzer, the data should display in the **Analysis Grid** viewer and in an interlaced fashion. Thereafter, you will create a **Union** that combines two fields of equal value but with different names into a single field with a new name, to simplify your data analysis processes with Message Analyzer.

IMPORTANT

Because creating a working **Union** in your Message Analyzer installation depends on combining fields that are specific to your environment, the procedure that follows uses hypothetical field names, such as `Command.smb_cmd` and `Command`. Therefore, when using this procedure to create a *working* **Union**, you should substitute actual field names that are contained in actual data files that are specific to your environment. In addition, you have the option to specify any **Union** name that is appropriate for your needs.

More Information

To learn more about the concepts upon which this example procedure is based, see Configuring and Managing Message Analyzer Unions.

To create a Union of two related data fields

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer.

   To obtain access to all features, be sure to run Message Analyzer as an Administrator.

2. From the Message Analyzer **Start Page**, click the **New Session** button to display the **New Session** dialog.

3. Under **Add Data Source** in the **New Session** dialog, click the **Files** button to display the **Files** tab along with the associated session configuration features that it contains in the **New Session** dialog for a Data Retrieval Session.

4. On the **Files** tab of the **New Session** dialog, click the **Add Files** button on the **Files** tab toolbar to launch the **Open** dialog, select the trace and log files that contain the data fields for which you will create a **Union**, and then click **Open**.

TIP

You will need to know in advance the field names from the messages in your input files for which you will be creating a new **Union**. As an example, the built-in **SMBTID** union that you can access from the global Message Analyzer **Tools** menu creates a **Union** of the following three fields that are accessible in **Field Chooser Tool Window** under the **SMB**, **SMB2**, and **SambaSysLog** nodes, respectively:

SMB.SmbHeader.Tid
SMB2.SMB2Request.Header.TreeId
SambaSysLog.smb_command.command.smb_tid - note that the **SambaSysLog** node and this field will only exist in **Field Chooser** if a Samba *.log file is loaded into Message Analyzer with a **SambaSysLog** configuration file specified in the **New Session** dialog for a Data Retrieval Session. This could apply if you are loading a *.log file into Message Analyzer.

Whenever you create a new **Union**, you will need to use the **Field Chooser** to locate and add the fields that are to comprise the **Union**, as indicated ahead in this procedure.

The name of the trace and log files appear in the files list.

5. Observe the current **Start With** drop-down list selection in the **New Session** dialog; if it is not the **Analysis Grid** viewer, click the drop-down list and select the **Analysis Grid** item.

6. Click the **Start** button in the **New Session** dialog and observe that the messages loaded into Message Analyzer from the log and trace files display in a chronological interlaced fashion in the **Analysis Grid** viewer, with a column for each differently named field of interest that displays similar values in each corresponding column.

7. Click the **Unions** button on the global Message Analyzer toolbar to display the **Edit Union** dialog.

8. In the **Edit Union** dialog, perform the following actions:

- In the **Name** text box, specify a name for the **Union**. Be sure to enter a name that is meaningful in your environment. In this example, the hypothetical **Union** name is `SMBCommand2`.
- In the **Category** combo-box, either select an existing **Category** or type a new one.
- To add the fields you want to combine in the **Union**, click the **Add** button to display the **Field Chooser Tool Window**, in which you can locate the field names. Note that you can add only one field at a time with the **Field Chooser**. In this example, the hypothetical field names are `Command.smb_cmd` and `Command`.

  As you add fields, you should notice the **Type** label displaying the most appropriate data type for the combined fields, as calculated by Message Analyzer; see Creating Unions for more information.
- When you are finished configuring the **Union**, click the **Save** button in the **Edit Union** dialog.

  The new **Union** is added to the root **Unions** node in the **Field Chooser** window.

9. Open the **Field Chooser** window by clicking the **Add Columns** button on the **Analysis Grid** viewer toolbar.
10. Expand the root **Unions** node in the **Field Chooser** and then double-click the name of the new **Union** to add it as a new data column in the **Analysis Grid** viewer column layout.

    Observe that the new <*unionName*> column in the **Analysis Grid** viewer correlates the data field values for the disparate field names that you specified in the **Union** you created. Note that you can remove the disparate field columns from the **Analysis Grid** viewer by selecting the **Remove** command that displays as a context menu item when you right-click the corresponding column header for each field. When the original field columns are removed from the **Analysis Grid** viewer, the **Union** name column will continue to correlate the values for the underlying data fields contained in the **Union**.

See Also

Configuring and Managing Message Analyzer Unions

Analyzing Message Data

23 minutes to read

Message Analyzer tools that are available for analyzing the results of a Live Trace Session or a Data Retrieval Session include data viewers, tool windows, and other features and functions that manipulate, locate, or otherwise interact with message data. These analysis tools are briefly summarized in this section.

Using the Session Analysis Tools
Message Analyzer provides many tools that you can use to achieve unique perspectives on your data set during an Analysis Session. The fundamental tool upon which all further analysis usually proceeds is the data viewer. By default, Message Analyzer provides the **Message Analyzer Chart View Layouts** asset collection Library that contains a diverse collection of built-in **Layouts** for the **Chart** viewer that you can select in any Analysis Session. These **Layouts** enable you to display data in diverse formats to expose analysis perspectives that can accelerate problem solving. Other tools that can manipulate or interact with the primary data display known as the **Analysis Grid** viewer include view **Filters**, **Viewpoints**, **Grouping**, column **Layouts**, **Aliases**, **Color Rules**, various **Tool Windows** such as **Details**, and so on.

A brief summary of the tools that you can use to perform data analysis tasks include the following:

- **Analysis Grid** viewer — the **Analysis Grid** is the primary analysis surface that presents message data in a tree grid format with rows of expandable message nodes in a stacked configuration. This common data viewer is usually the starting point for most Analysis Sessions, as it enables you to review a core set of data for any trace, which includes default column data and message summaries. In addition, selection of messages in the **Analysis Grid** viewer interactively drives the display of related details that appear in the **Message Stack**, message **Details**, **Message Data**, and **Field** value windows. The **Analysis Grid** viewer also provides access to various data manipulation features from a right-click context menu, such as **Filters**, **Pattern Expressions**, and **Time Shifts**, in addition to numerous toolbar functions.

More Information
To learn more about the **Analysis Grid**, see the Analysis Grid Viewer topic.

- **Chart** viewer **Layouts** — enable you to display data in different graphical formats, such as **Bar** element, **Pie** chart, **Table** grid, or **Timeline** graph, to create unique analysis perspectives. The **Layouts** for the **Chart** viewer include graphic visualizer components that can provide high-level data summaries or low-level data that is focused on specific message types and values. These include **Layouts** such as the **IP/Ethernet Conversations by Message Count**, **Average Response Times for Operations**, **SMB Top Commands**, and so on, in addition to any custom **Layouts** that you create. All of the built-in **Layouts** for the **Chart** viewer are designed to provide a unique focus on your data in order to expose different values and details through graphic visualizer components, which can be critical when you need to expedite the data analysis process.

More Information
To learn more about the built-in **Layouts** for the **Chart** viewer, see the Chart Viewer Layouts topic.

- **Grouping** viewer — enables you to organize the traffic from a set of trace results into summary hierarchies that are based on built-in or user-customized **Grouping** viewer **Layouts**, which contain nested Groups that are defined by various fields, properties, annotations, methods, and so on. You can alter the nesting configuration of the Groups for different analysis perspectives and you can correlate message

selection and filtering on an interactive basis with the **Analysis Grid** viewer.

For example, selection of groups in the **Grouping** viewer interactively drives the display of messages in other viewers of the same session, which enables you to create analysis contexts that focus on specific groups of messages in different presentation formats. With the **Grouping** viewer, you can drill down into the data of lower-level Groups in the nested group configuration to isolate and expose discrete information based on the type of Group selected. The **Grouping** viewer is a primary analysis tool that can help you quickly achieve the following:

- Organize data into unique hierarchies to expose targeted information from large data sets.
- Identify Groups that reflect the highest traffic volumes.
- Isolate all messages in a set of trace results to a specific top-level Group and drill down for data in nested Groups, to obtain a concise analytical focus at each Group level.
- Correlate messages in the **Analysis Grid** and other viewers with the Groups in which they appear in the **Grouping** viewer.
- Correlate messages across different data sources, such as a log and a saved trace file.

More Information
To learn more about the **Grouping** viewer, see the Grouping Viewer topic.

- **Pattern Match** viewer — enables you to select a **Pattern** expression from a Library of predefined **Patterns** from Microsoft, that can aggregate a collection of messages from a set of trace results that all contain a common pattern or sequence that matches the criteria of the applied **Pattern** expression. A pattern might consist of a sequence of fields, properties, values, annotations, or other relationships. This feature can help you expose the context or sequence in which certain events occurred across the timeline of a trace session, as opposed to filtering, where the impact is typically limited to the boundary of individual messages rather than to groups of messages. The **Pattern Match** viewer contains an **AVAILABLE PATTERNS** Library from where you can select predefined **Pattern** expressions that provide pattern locating functions that can be useful in most environments. The **Pattern Match** viewer also provides a **Create Pattern** button that opens the **Pattern Editor** dialog from where you can create your own **Pattern** expressions.

  When you click the **Create Pattern** button in the **Pattern Match** viewer to open the **Pattern Editor** dialog, it opens to the **Quick** tab, which contains an **Insert Message** button that enables you to select message fields from the **Field Chooser Tool Window**, after which the **Quick** tab will be populated with initial data for a message type that you selected, along with other UI automation that you can access by clicking the **Insert Criteria** link. You will also notice a **Free Form** tab in the **Pattern Editor** that contains an editing surface where you can write a **Pattern** expression in Open Protocol Notation (OPN) code without the support of any UI automation, although you would need to be familiar with OPN to do so.

  Note that you can also open the **Pattern Editor** dialog by right-clicking a message in the **Analysis Grid** viewer and then selecting the **Create Pattern** command in the context menu that appears. However, in this case, the **Pattern Editor** dialog opens to the **Quick** tab with UI automation controls provided and seed data already inserted, to help you develop a **Pattern** expression more quickly. Moreover, the **Quick** tab is pre-populated with initial information that is derived from the message/s you selected in the **Analysis Grid** viewer, the assumption being that you want to create a **Pattern** expression based on a sequence of fields, properties, or values (in some relationship to one another) from the selected message type, and which you will define in the expression you create.

More Information
To learn more about using the **Pattern Match** viewer and creating your own **Pattern** expressions, see the Pattern Match Viewer section.

- View **Filters** — enable you to apply a built-in or custom Filter Expression to a set of trace results to isolate specific messages that can expose errors, values, or other details, in order to pinpoint the cause of a particular problem that is occurring. Only the messages that pass the criteria of the applied filter are returned, while all others are temporarily removed from the results, so you can create a focused set of messages for analysis. A view **Filter** is a handy analysis tool because you can apply and remove them without impacting the original set of trace results, or you can specify a different view **Filter** to create a new set of results that present a different analysis perspective. You can select built-in view **Filters** from a centralized **Library** or configure custom Filter Expressions of your own on any Filter panel text box on the Filtering toolbar just above the data viewing surface.

More Information
To learn more about view **Filters**, see Applying and Managing Filters.

- **Viewpoints** — enable you examine network traffic from the perspective of a particular protocol or stack layer, with no messages above it. By applying a built-in **Viewpoint**, you can remove messages above the **Viewpoint** to create a focused set of messages that can expose details that are normally hidden within the message stack. A **Viewpoint** can create a unique analysis perspective on your data, while at the same time streamlining the analysis process. For example, if you want to view and troubleshoot TCP messages only, you can select the **TCP** viewpoint to drive TCP messages, including fragments, to top-level in the **Analysis Grid** viewer, with all upper layer protocols or modules removed from view. Message Analyzer provides numerous built-in **Viewpoint** configurations that are contained in a Library that is accessible from the **Viewpoints** drop-down list on the Filtering toolbar just above the data viewing surface.

  This drop-down list also has a **Viewpoint** called **Disable Operations** that enables you to break up all the request/response message pairs that are encapsulated in top-level Operation nodes in the **Analysis Grid** viewer, for protocols that make use of these pairs, such as HTTP, DNS, SMB, and LDAP. This reestablishes the chronological context in which messages were originally captured, similar to the Network Monitor view, although this makes response messages more difficult to locate and correlate with their associated request messages. In addition, you have the option to apply a **Viewpoint Filter** that is accessible from the centralized **Library** on the Filter panel that displays when you select the **Add Viewpoint Filter** command from the **Add Filter** drop-down list on the Filtering toolbar. This type of filter enables you to further refine the focus of messages by applying filtering *after* a **Viewpoint** is applied. However, note that you can configure and apply a **Viewpoint Filter** only after a **Viewpoint** is applied to a set of trace results.

More Information
To learn more about **Viewpoints**, see Applying and Managing Viewpoints.

- **Parsing** options — enable you to reparse trace results based on an alternate port specification that differs from the standard port typically used by a particular protocol. It is becoming a common practice to pass traffic to alternate ports to avoid exposing network traffic to exploitation. The **Parsing** options provide a range of protocols for which you can specify an alternate port to accommodate traffic that used such an alternate port. For example, you might specify alternate port 8080 instead of HTTP standard port 80, or alternate port 3268 for LDAP standard port 389.

  The protocols for which you can specify alternate ports are contained in a drop-down list on the **Parsing** tab of the **Options** dialog, which is accessible from the global Message Analyzer **Tools** menu. You can also launch this dialog by right-clicking a message in the **Analysis Grid** viewer and then selecting the **Parse As** command from the context menu that appears. When you specify one or more alternate port numbers for a particular protocol, the OPN parser for that protocol uses the port number/s that you specified when your trace is reparsed, so you can view the messages received on the specified ports.

More Information
To learn more about **Parsing** options, see the **Parsing** section of Setting Message Analyzer Global Options.

- **Shift Time** dialog — enables you to change the **Timestamp** of captured messages that are displayed in an Analysis Session. The time shift capability enables you to compensate for skewed system clock values or time zone differences across different computers, so that you can chronologically align the messages from those computers and ensure accurate troubleshooting.

More Information
To learn more about the **Shift Time** feature, see Setting Time Shifts.

- **Aliases** — enables you to substitute more friendly names for several types of data field values in the **Analysis Grid** viewer, to facilitate easier recognition of values that can otherwise be cryptic and difficult to work with, for example IPv4 and particularly IPv6 addresses. This feature can improve your ability to discover and analyze specific message traffic through the use of simplified names that have meaning in your troubleshooting environment. By customizing your data analysis environment with **Aliases**, keeping track of traffic to and from host IP addresses, physical addresses, and ports becomes easier. By default, Message Analyzer provides two **Loopback Aliases** that you can select from the **Aliases** drop-down list on the Message Analyzer global toolbar.

More Information
To learn more about the **Aliases** feature, see Using and Managing Message Analyzer Aliases.

- **Color Rules** — an important analysis feature for the **Analysis Grid** viewer that enables you to use color, text, and font styles to decorate and highlight messages that contain specific types of information that you can identify at-a-glance, which can preclude the need for additional analysis and diagnostics. **Color Rules** can provide an instant visual cue of messages that meet predefined criteria, to alert you that closer scrutiny and further investigation may be needed. By default, Message Analyzer provides a **Color Rule** Library that contains a host of built-in **Color Rules** that you can use immediately. You can also create your own **Color Rules** to meet the needs of your particular environment.

More Information
To learn more about **Color Rules**, see Using and Managing Color Rules.

- **Go To Message** — an important analysis feature for the **Analysis Grid** viewer that enables you to quickly locate any message by number in a large collection of messages. If you are working with multiple data sources, the **Go To Message** dialog provides options to search across all data sources for a particular message number that you specify, or you can select a specific data source to search. By using this feature, you can accelerate your analysis process, especially in support scenarios where you are being directed to find particular messages very quickly.

More Information
To learn more about the **Go To Message** features, see Using the Go To Message Feature.

- **Layouts** — an essential analysis feature for the **Analysis Grid** viewer that enables you to select view **Layout** configurations that are tailored for different types of message analysis. View **Layouts** add new data columns to the **Analysis Grid** to expose typical field information of specific protocols that is related to issues that are commonly investigated for such protocols. For example, if you are troubleshooting TCP, you might select one of the **TCP Deep Packet Analysis** view **Layouts** to populate your **Analysis Grid** viewer with data that enhances this type of analysis.

  Message Analyzer provides numerous view **Layouts** that are customized for common analysis scenarios. You can access these from the **Layout** drop-down list on the **Analysis Grid** toolbar. You can also create your own custom view **Layouts** by adding columns to the **Analysis Grid** viewer with the **Field Chooser Tool Window** and then saving your configuration with the **Save Current Layout As...** command from

the **Layout** drop-down list.

> **NOTE**
>
> Message Analyzer also provides the **Profiles** feature, which displays preset configurations of data viewers and **Layouts** that are specifically designed to create a unified and interactive analysis environment that is tailored for the type of data you are loading into Message Analyzer. For example, if you are loading data from .cap files, you can select a **Profile** such as the **Performance Top Down Profile**, that will display the **Performance Top Down** layout for the **Analysis Grid** viewer, the **Process Names and Conversations** layout for the **Grouping** viewer, and the **Top TCP/UDP Conversations by Message Count** layout for the **Chart** viewer, providing that you enable the **Performance Top Down Profile** prior to loading data from such an input file.
>
> For further details about Message Analyzer **Profiles**, see Working With Message Analyzer Profiles.

More Information
To learn more about **Analysis Grid** viewer **Layouts**, see [Applying and Managing Analysis Grid Viewer Layouts](applying-and-managing-analysis-grid-viewer-layouts.md).
To learn more about **Grouping** viewer **Layouts**, see the Grouping Viewer topic.
To learn more about **Layouts** for the **Chart** viewer, see the Chart Viewer Layouts topic.

- **Unions** — an important analysis feature for the **Analysis Grid** viewer that enables you to correlate data from varying field names in multiple data sources, where the data is of the same type but the field names are different. This feature can streamline the analysis of logs and saved traces from a common environment, where the file formats of these message sources use different naming conventions to identify the same data fields. When this is the case, Message Analyzer can display this field data from multiple sources in a single **Analysis Grid** data column that is named by a **Union** that you create.

  By default, Message Analyzer provides several common **Unions** that exist in the **Unions** node of the **Field Chooser Tool Window**. You can add one or more of these **Unions** as new columns in the **Analysis Grid**; however, it will be useful only if it is appropriate for the data you are analyzing. For example, if you are working with an SMB trace and a SambaSysLog, you might add the **SMBCommand** union as a new column to merge the **SMB.Command** field of the trace and the **SambaSysLog.smb_command.command.smb_com** field of the log. You can add this new column to the **Analysis Grid** viewer by right-clicking the **Union** and selecting the **Add as Column** command in the context menu that appears.

  You can assess the data fields that are merged with any of the built-in **Unions** by opening the **Manage Unions** dialog that is accessible from the **Unions** drop-down list in the global Message Analyzer **Tools** menu. In the **Manage Unions** dialog, right-click a **Union** in the list and select the **Create a Copy** command in the context menu that appears, which in turn displays the **Edit Union** dialog. The fields that are merged by the selected **Union** appear in the **Select fields to include** list box of the **Edit Union** dialog. Note that you can also manually configure your own **Unions** by creating a single field entity that correlates multiple disparate field names from different data sources with the use of the **Edit Union** dialog, which is accessible as specified immediately above, or by clicking the **Unions** button on the global Message Analyzer toolbar.

More Information
To learn more about **Unions**, see Configuring and Managing Message Analyzer Unions.

- **Diagnostics** — a feature that can be critical when attempting to assess the cause of errors and failures in a set of trace results. Message Analyzer provides two different tools that provide diagnostic information. These tools are similar but they present data in different ways. The first is the **DiagnosisTypes** column in the **Analysis Grid** viewer. If an error occurred for a particular message, one of four different icons appear in the **DiagnosisTypes** column for that message; however, the diagnostic information can be scattered

across the trace or buried deep in the message stack, making it a little harder to analyze.

One technique you can use to gather the messages that have diagnostic information is to sort the **DiagnosisTypes** column. But this does not necessarily expose the actual messages that contain a diagnosis condition, since Message Analyzer bubbles up the icons to top-level as a cue that you have an embedded diagnosis message. In addition, by simply sorting, it might not be readily apparent as to what the diagnostic message descriptions indicate nor does it expose how many contain the same message description, without performing a laborious manual process of repetitive node expansions in the **Analysis Grid** viewer. To alleviate this problem, Message Analyzer provides a **Diagnostics Tool Window**, which contains tabular data that summarizes the diagnosis type, the module in which diagnosis messages occurred, the message descriptions, and the number of identical messages for each diagnosis type. In addition, selection of a diagnosis row in the **Diagnostics Tool Window** table interactively drives message selection in the **Analysis Grid** viewer for immediate data correlation. At a single glance, these summaries can provide an instant assessment of where problems exist in the context of an entire trace.

More Information
To learn more about Message Analyzer diagnostics, see the Diagnostics Tool Window and Diagnosis Types topics.

- **Message selection** — Message Analyzer provides some very useful enhancements to the message selection process, by enabling you to track the messages that you select in any particular part of an Analysis Session. This capability is provided by the **Selection Tool Window**, which is accessible from the **Windows** submenu of the Message Analyzer global **Tools** menu. By building a collection based on your message selections, Message Analyzer enables you to backtrack to specific messages that you selected earlier during analysis, to revisit them for further scrutiny. You can also forward track through the selection collection for the same purpose. This feature can be of benefit if you are analyzing multiple messages in a large set of data. It is also advantageous if you accidentally lose focus on a message, because you can return to any previous selection with a single click on the **Go back** or **Go Forward** button on the **Selection** window toolbar.

  You can also turn message **Selection** tracking on and off, at your discretion. Also note that message selection is interactive with other viewers and some **Tool Windows** for enhanced correlation capabilities. You can even navigate message selections that you made across different Message Analyzer sessions.

More Information
To learn more about message selection capabilities, see the Selection Tool Window topic.

- **Tool Windows** — the Message Analyzer **Tool Windows** provide a vast amount of functionality that enhances the data analysis process. Some **Tool Windows** interact with data viewers, for example, message selection in a **Tool Window** may drive message highlighting in a data viewer such as the **Analysis Grid**, and vice versa. Moreover, message selection in the **Analysis Grid** can also drive the display of various message details in **Tool Windows** such as the **Message Stack**, **Details**, and **Field** windows, which provide a core set of analysis tools that you will no doubt use on a consistent basis. Some **Tool Windows** are session-specific, for example the **Diagnostics** window, which means that they respond to session selection and are not driven by message selection in other tools or viewers. Some **Tool Windows** are message-specific, for example the **Details** window, which means that they do respond to message selection in other tools or viewers.

More Information
To learn more about the **Tool Windows** that Message Analyzer provides, see the Tool Windows section.

- **Field Chooser** — this **Tool Window** enables you to enhance the functionality of other analysis tools. It displays the message hierarchy for all modules and protocols for which Message Analyzer contains OPN parsers. The **Field Chooser** message hierarchy consists of message types, fields, and properties for all

modules with expandable nodes in a tree format. It also contains various global fields, annotations, and properties that you can utilize in various ways. The following are a few examples of how you can use the **Field Chooser**:

- Add message fields as columns to the **Analysis Grid** viewer and **Grouping** viewer, to enhance the scope of the data that these viewers present.
- Select fields or properties that you want to use in a data formula for a new **Chart** viewer **Layout** you are creating.
- Select fields or properties that you want to use in a **Pattern** expression you are creating.
- Correlate disparate field names of the same type from different data sources into a single entity when creating a new **Union**.
- Review message hierarchies to discover field and property names of various protocols, to use when you are creating Filter Expressions.

More Information
To learn more about the **Field Chooser**, see the Using the Field Chooser and Field Chooser Tool Window topics.

- **Analysis Grid Group** command — a handy feature that enables you to organize trace data into groups based on the varying values of an **Analysis Grid** viewer data column, which is named by and displays the values of a field, property, annotation, or method across a set of trace results or log content. By organizing values into groups, you can quickly isolate messages of interest and make it easier to discover data for which you are searching. For example, if you right-click the **DiagnosisType** column in the **Analysis Grid** viewer and select the **Group** command, you can isolate your data into groups with expandable nodes where each group contains the same type of diagnosis message. In addition, you can perform subsequent **Group** commands in like manner to create a *nested* group configuration that enables you to correlate data among the groups. Lastly, you have the option to reorder the nesting configuration by dragging and dropping group labels into different positions of the initial group order. Reorganizing your trace data into groups or nested group configurations with the **Group** command enables you to create unique analysis perspectives that facilitate rapid identification of issues.

More Information
To learn more about the **Analysis Grid** grouping feature, see Using the Analysis Grid Group Feature.

- **Analysis Grid column filtering** — a feature that enables you to conveniently isolate messages according to a specified search string that matches a value in a particular data column. You simply click the **Show Column Filter Row** icon on the left side of the column label row to display an amber-colored search box below each column label. Thereafter, you can enter a search string and Message Analyzer automatically reorganizes the data display to include only the messages that contain a column value that matches your search string, that is, if any such values are found.

  For example, if you typed the name of a protocol in the amber text box below the **Module** column label, such as HTTP, then Message Analyzer will temporarily remove all other messages in the trace results except HTTP. When you remove the search string, Message Analyzer reinstates the original message set. You also have the option to specify another search string as a column filter to further narrow down your results. For instance, after specifying a column filter such as HTTP in the previous example, you might also specify a search string for the **Summary** column with the text "GET", to filter out all messages from the display, except HTTP requests that use the GET method. This feature can help you quickly locate specific messages with specific field values in a large data set and reduce your analysis time.

More Information
To learn more about the column filtering feature, see the Filtering Column Data topic.

- **View**
  Field and Property values — a core analysis feature that enables you to view the values of any field or property in messages that were parsed by Message Analyzer. These values will display in the **Details Tool Window**, which appear when you select a message in the **Analysis Grid** viewer. The **Details** window typically shows the field or property **Name**, **Value**, **Bit Offset**, **Bit Length**, and **Type**. It is likely that you will use this **Tool Window** extensively in analysis, because it immediately exposes the values of any message field.

More Information
To learn more about viewing message field and property values, see the Message Details Tool Window topic.

- **Track Field and Property values** — this feature adds a new dimension to field and property analysis, by enabling you to select specific fields or properties for which you want to track values across a message set. You will find this capability on the toolbar of the **Details Tool Window**, which displays the following icons. You can identify these icons by their hover-over tool tips:
  - **Show all fields for the selected message** — click this icon to show all fields for a message that you selected.
  - **Show all properties for the selected message** — click this icon to show all properties that apply to a message that you selected.
  - **Show tracked fields and properties for the selected message** — displays a list of fields and/or properties that you have set for tracking.

  You can set a field or property for tracking by right-clicking it in the **Details** window and then selecting the **Track 'field/propertyName'** command, where 'field/propertyName' is a placeholder for an actual field or property name. Thereafter, you can display the tracking list and then either scroll through or arbitrarily select **Analysis Grid** messages to view tracked values. You might also select messages in the **Grouping** viewer and monitor tracked fields or properties in **Details**. Moreover, if the **Selection Tool Window** is enabled for selection tracking, you can forward- and back-navigate through your selections to assist your analysis process.

More Information
To learn more about tracking message field and property values in the **Details** window, see the Using the Details Tool Window Features topic.

Filtering Message Data

2 minutes to read

This section describes various ways to filter message data in Message Analyzer so that you can focus on traffic that is isolated by specific message type, field value, address, port communication, string value, diagnosis type, and so on. You can use Filter Expressions to select specific data from saved files or live captures when a Data Retrieval Session or Live Trace Session is running, respectively. You can also apply Filter Expressions when you are working with trace *results* to narrow down your view to specific data of interest while filtering out all the rest. In addition, you can either specify built-in Filter Expressions, or you can manually create your own.

You can even create a Filter Expression that uses an **Alias** (typically a friendly name that replaces some cryptic field value; see Using and Managing Message Analyzer Aliases), and you can save such a filter in the centralized Filter Expression **Library**. This enables you to use the Filter Expression that incorporates an **Alias** as a **Session Filter** or a view **Filter**. Note that you can also create Filter Expressions that include **Unions**.

What You Will Learn

In the topics of this section, you will specifically learn about Message Analyzer filtering capabilities, which includes filtering imported data in a Data Retrieval Session, filtering data while it is captured in a Live Trace Session, filtering trace results, and learning how to create your own Filter Expressions, as described below.

In This Section

Filtering Loaded Input Data — apply a built-in or user-developed **Session Filter** or a **Time Filter** to a Data Retrieval Session to constrain input message volume; also use input file selection as a method of limiting input data.

Filtering Captured Input Data — apply numerous types of driver-level filters to focus on specific types of messages in a trace, so you can limit the amount of data you will capture in a Live Trace Session for problem solving. Provider-level filters include **Fast Filters**, **Keyword** bitmask filters, **Level** filters, **WFP Layer Set** filters, **HTTP** filters, host adapter NDIS layer filters, and Hyper-V-Switch extension layer filters.

Filtering Live Trace Session Results — assess the functions of built-in view **Filters**, so you can better utilize them to apply filtering to a set of trace results.

Filtering Column Data — make use of the **Analysis Grid** viewer **Column Filter** feature to quickly display only messages with fields that contain search text that you specify in a **Column Filter Row** text box.

Writing Filter Expressions — write your own Filter Expressions with the Message Analyzer Filtering Language, which is based upon the Open Protocol Notation (OPN) language.

Go To Procedures

To proceed directly to procedures that demonstrate the filtering features described in this task area, see **Procedures: Using the Data Filtering Features**.

See Also

Applying and Managing Filters

Filtering Loaded Input Data

6 minutes to read

Prior to loading data into Message Analyzer through a Data Retrieval Session, there are several methods you can use to limit the data that you load for analysis. These methods, which enable you to create a focused message set based on criteria that you define, consist of the following:

- **Configure a Session Filter** — enables you to "select" specific data from your input file configuration through filtering, to narrow the scope of data retrieval to specific data of interest, reduce the volume of messages that you load, and therefore improve performance.
- **Choose input files** — enables you to select data by combining chosen input files that contain specific data that you want to load into Message Analyzer.
- **Apply a Time Filter** — enables you to select data by reducing the scope of loaded messages to a specified time range, by configuring a window of time in which to view data prior to loading such data from saved files.
- **Choose a Parsing Level** — enables you to select data by limiting how far up the message stack Message Analyzer will parse, thereby reducing the number of messages processed and displayed, creating a focused analysis level, and increasing performance.

Selecting Data with a Session Filter

You can specify a **Session Filter** for your Data Retrieval Session by either writing your own Filter Expression or choosing a built-in Filter Expression from the centralized filter **Library**. This **Library** is accessible in the **New Session** dialog on the toolbar just above the **Session Filter** text box. This **Library** contains the same built-in Filter Expressions that you will find in the **Message Analyzer Filters** asset collection **Library** on the Filtering toolbar that appears just above the main analysis surface in the Message Analyzer user interface (UI).

If you want to specify a built-in Filter Expression from the centralized filter **Library** when loading data into Message Analyzer, you should first review the filter functions that are described in Filtering Live Trace Session Results and then select an appropriate **Session Filter**. If you intend to create your own Filter Expression, see the appropriate topic below in the **See Also** section for more information about how to write one.

> **TIP**
> After you write a Filter Expression or modify a built-in expression that you want to use as a **Session Filter**, you have the option to save the new Filter Expression to the centralized filter **Library** from the **New Session** dialog. For example, if you want to save the Filter Expression code that you specified in the **Session Filter** text box, click the **Library** drop-down list and then select the **New Filter** command to display the **Edit Filter** dialog. The text of the target Filter Expression is transferred to the **Edit Filter** dialog, where you can specify **Name**, **Description**, and **Category** information before you save it to the **Library** (by clicking the **Save** button in the dialog).

Saving a custom Filter Expression to the filter **Library** is optional. Because you are not restricted from applying such a filter to the data loading process, you can observe its performance prior to saving it. To do this, simply click the **Start** button in the **New Session** dialog.

> **NOTE**
> Message Analyzer no longer automatically validates that Filter Expressions properly compile, in order to remove the restrictions on Filter Expressions that use elements which are outside the range of the modules that Message Analyzer parses.

If your Filter Expression performs as expected and proves to be a useful tool for analysis, you can save it to your centralized filter **Library** for future use, as previously described. Note that such a Filter Expression is added to the **My Items** category of your Filter Expression **Library** and is then available to the Sharing Infrastructure, which enables you to share filters with other users.

More Information
To learn more about using a **Session Filter**, see Applying a Session Filter to a Data Retrieval Session.

Selecting Data by Choosing Files

After using the **Add Files** command on the toolbar of the **Files** tab in the **New Session** dialog to create a list of files that contain the data you want to load through your Data Retrieval Session, feasibly from multiple disparate data sources, you can also apply filtering through *file selection* to create unique message collections. To do this, simply select specific files from among the files that are marked for loading data in the files list to create a subset of messages that you want to focus on. For each file that you want to include for data loading, place a check mark in the check box next to the file name in the files list of the **New Session** dialog. You can also specify a data viewer in which to present the loaded data by selecting one from the **Start With** drop-down list. After you click the **Start** button in the **New Session** dialog, the data loading process is silently invoked in the background.

Selecting Data Through Time Filtering

You might have multiple log files containing data that was collected over a period of time from different sources, for example, from a client and server. For example, you might be interested in isolating the source of a TCP connection issue that involved lost TCP segments but it is unclear which machine was dropping packets. You could load data from sets of log files from the same time period, possibly while drilling down to a specific window of time with the use of an input **Time Filter**, to select data that was collected on both client and server computers in a specified time window. When viewing results in an Analysis Session, the data from each set is presented as a single unified message collection with messages interleaved in chronological order. Note that you might also want to configure a **Time Shift** if the data sets need to be synchronized.

More Information
To learn more about creating a **Time Filter** window, see Applying an Input Time Filter to a Data Retrieval Session.
To learn more about creating a **Time Shift**, see Setting Time Shifts.

Selecting Data by Choosing a Parsing Level

Message Analyzer provides a built-in Library known as the **Message Analyzer Parsing Levels** asset collection, as indicated in the **Asset Manager** dialog which is accessible from the global Message Analyzer **Tools** menu. This built-in Library is accessible from the **Parsing Level** drop-down list in the **New Session** dialog. In the list, the **Full** value is the default, which means Message Analyzer will parse the full stack for all messages. However, as an example, if you select the **Network Analysis** item in the list, Message Analyzer will parse up to and including the IP/Network Layer, with a few exceptions for certain TCP, UDP, and other traffic that is predetermined to be valuable for network analysis. Other **Parsing Levels** also use exceptions that are valuable to analysis at the chosen **Parsing Level**.

More Information
To learn more about **Parsing Levels**, see Setting the Session Parsing Level.

Changing the Applied Filtering by Editing a Session

You also have the option to edit your Data Retrieval Session after the data loads into Message Analyzer. You can do this by clicking the **Edit Session** button on the global Message Analyzer toolbar. Thereafter, the **Edit Session** dialog displays, from where you can choose another **Session Filter** or configure a different one, that is, if you exit the **Restricted Edit** mode. This includes reconfiguring any specified input **Time Filter** and/or choosing a different **Parsing Level** as part of your edits. You can then apply the changes you make to the session configuration by clicking the **Apply** button in the **Edit Session** dialog and Message Analyzer will process the changes you specified for your Data Retrieval Session. Note that you can add or remove files in the files list on the **Files** tab without exiting the **Restricted Edit** mode, however, for any other type of edit, you must click the **Full Edit** button to enable all editable features.

More Information
To learn more about editing a Data Retrieval Session, see Editing Existing Sessions.

See Also

Selecting Data to Retrieve
Selecting Data to Capture
Writing Filter Expressions

Filtering Captured Input Data

6 minutes to read

Prior to starting a Live Trace Session with Message Analyzer, there are numerous types of filter configurations that you can create to limit the scope of data that you capture. An overview of the filter types that you can apply to a Live Trace Session is provided in this section.

Using a Session Filter

One of the simplest and most effective methods of filtering is to add a **Session Filter**. A **Session Filter** will allow you to retrieve only the messages that meet the filtering criteria that you define. This conveniently provides a way to target specific message data while reducing the number of retrieved messages for better performance. Just as you can do in a Data Retrieval Session, you can either select a built-in **Session Filter** or configure your own in the **Session Filter** text box of the **New Session** dialog. If you want to use a built-in **Session Filter**, you can select one from the **Message Analyzer Filters** asset collection **Library** that appears on the toolbar above the **Session Filter** text box in the **New Session** dialog, or you can create your own by entering filter parameters in the indicated text box during Live Trace Session configuration. For additional information about **Session Filters**, see Working with Session Filters in a Live Trace Session or Applying a Session Filter to a Data Retrieval Session.

Using Other Filters

In addition to specifying a **Session Filter** when you are configuring a Live Trace Session, you can also specify any of the following filter types:

- **Fast Filters** — if you are configuring a Live Trace Session with a **Trace Scenario** that uses the **Microsoft-PEF-NDIS-PacketCapture** provider, you have the option to specify up to three **Fast Filters** that operate efficiently at the kernel level. You can also specify up to four **Fast Filters** when configuring a Live Trace Session with any **Trace Scenario** that uses the **Microsoft-PEF-WFP-MessageProvider**.

  A **Fast Filter** is a capture-mode filter that offers a significant improvement in performance over user-mode filtering, for example, when applying a **Session Filter**. In the former case, capture-mode filtering is instrumented at the driver level before messages are delivered to the PEF Runtime, while in the latter case, user-mode filtering is applied as part of the Runtime parsing process. User-mode filtering therefore adds more processing time before Message Analyzer can access the data from the PEF Runtime API for display.

- **Network Adapters and Fast Filter Groups** — if you are configuring a Live Trace Session with a **Trace Scenario** that uses the **Microsoft-PEF-NDIS-PacketCapture** provider on computers running the Windows 7, Windows 8, or Windows Server 2012 operating system, you can also specify the network adapters through which your Live Trace Session will capture messages. For example, you can isolate messages to an Ethernet adapter, a wireless adapter, a combination of adapters, and so on. In addition, you can assign **Groups** of **Fast Filters** to any local adapter.

  When you install Message Analyzer, all the network adapters on your system are enumerated. If you open the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog to configure filters for the **Microsoft-PEF-NDIS-PacketCapture** provider in the **Local Network Interfaces Trace Scenario**, you will see that all the network adapters on your system are populated to the **System Network** tree grid under the **Machine** node in the **Advanced Settings** dialog for the **Microsoft-PEF-NDIS-PacketCapture** provider. This dialog contains the configuration features that enable you to create filter **Groups** and assign them to specific adapters.

  In addition, you can selectively enable or disable any network adapter that appears in the **System Network**

tree grid of the dialog. By isolating the network adapter on which you capture data, you can block messages from other adapters and focus on capturing the messages of a particular protocol, for example, the Point-to-Point over Ethernet (PpoE) protocol on a WAN Miniport interface connection. Also, you can create filtering configurations that apply to all adapters, a group of selected adapters, or a single selected adapter only.

The features of the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog provide a flexible framework that enables you to focus on capturing very specific data while achieving the performance advantages that are inherent to **Fast Filters**, as described in Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

- **NDIS Layer and Hyper-V-Switch Extension Filtering** — if you are configuring a **Remote Network Interfaces Trace Scenario** on computers running the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system, which uses the **Microsoft-Windows-NDIS-PacketCapture** provider with remote capabilities, you can specify how packets are intercepted on the NDIS filter layers of a remote host adapter or on the extension layers of a Hyper-V-Switch that services virtual machines (VMs) on which you are monitoring traffic. You can also specify the direction that packets traverse these layers along with other special filter configurations that specify a **Truncation** value, **EtherTypes**, **IP Protocol Numbers**, **MAC addresses**, and **IP addresses**. You can also specify particular remote host or VM adapters on which to capture data while excluding others. The configuration for such settings is available in the **Advanced Settings - Microsoft-Windows-NDIS-PacketCapture** dialog, which is described in detail in Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

- **WFP Layer Set filters** — if you are configuring a Live Trace Session with any **Trace Scenario** that uses the **Microsoft-PEF-WFP-MessageProvider**, you can specify **WFP Layer Set** filters that isolate IPv4 or IPv6 message traffic directionally at the Transport layer. The **WFP Layer Set** consists of kernel-mode TCP/IP stack filters that operate in the receive or send path at the Transport layer. These filters allow you to selectively enable or disable either all inbound or all outbound packets at the Transport layer when capturing IPv4 or IPv6 messages.

- **HTTP filters** — if you are configuring a Live Trace Session with the **Pre-Encryption for HTTPS Trace Scenario** that uses the **Microsoft-Pef-WebProxy** provider, you can specify filters that isolate traffic based on a **Hostname** or **PortFilter** value.

- **Keyword and Level filters** — if you are configuring a Live Trace Session that uses a particular system **ETW Provider**, you can set event **Keyword** bitmask and **Level** filters to capture events from specific modules of a Windows system component that has been instrumented for ETW via that **ETW Provider**, with **Keyword** bitmask values and **Level** strings that represent its events. By setting an appropriate **Keyword** bitmask or **Level** value, you cause the **ETW Provider** to deliver only the events that are represented by the **Keyword** or **Level** configuration, thereby enabling you to filter for these events in traces that use a particular system **ETW Provider**. Examples of such providers include the **Microsoft-Windows-Dhcp-Client** and **Microsoft-Windows-LDAP-Client**. These ETW providers are accessible from the **Add System Providers** dialog, which you can display from the **Add Providers** drop-down list on the **ETW Providers** toolbar in the **New Session** dialog during Live Trace Session configuration.

> **NOTE**
>
> The default **Microsoft-PEF-NDIS-PacketCapture** provider and **Microsoft-PEF-WFP-MessageProvider** both enable you to specify event **Keyword** bitmask and **Level** filter configurations. However, you should note that not all system **ETW Providers** are enabled for event **Keyword** and **Level** configuration, as some **ETW Providers** do not define them.

More Information

To learn more about the functions of built-in Filter Expressions that you can apply as a **Session Filter** or view **Filter**, see Filtering Live Trace Session Results.

To learn more about creating your own Filter Expressions or modifying existing ones, see Writing Filter Expressions.
To learn more about the **Fast Filter** configurations that you can apply to PEF providers, see Common Provider Configuration Settings Summary.
To learn more about **Network Adapter** filtering, see Common Provider Configuration Settings Summary.
To learn more about **Fast Filter Groups** for adapters, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.
To learn more about remote tracing and special NDIS layer and Hyper-V-Switch extension filters, see Configuring a Remote Capture.
To learn more about **WFP Layer Set** filtering, see Common Provider Configuration Settings Summary.
To learn more about **Hostname** and **PortFilter** filtering, see WebProxy Filters.
To learn more about event **Keyword** bitmask and **Level** filtering, see System ETW Provider Event Keyword/Level Settings.

See Also

Configuring a Live Trace Session

Filtering Live Trace Session Results

17 minutes to read

After you capture data in a Live Trace Session, load data through a Data Retrieval Session, or load data using the **Open** or **Recent Files** features, you will typically analyze your trace results in one of the Message Analyzer data viewers, such as the default **Analysis Grid** viewer. As part of analysis, it is likely that you will need to manipulate the trace results to expose specific information that you want to examine. In Message Analyzer, a primary method for isolating specific data during trace results analysis is to apply a view **Filter**, which temporarily alters the data that Message Analyzer displays according to applied filtering criteria. When you finish analyzing the effects of the applied view **Filter**, you can simply click the **Remove** button on the Filtering toolbar to remove the filter and return to the original trace session results.

You can specify a view **Filter** in any of the following ways:

- **Predefined configuration** — select a built-in Filter Expression from the **Message Analyzer Filters** asset collection in the centralized user **Library** on the Filtering toolbar, which is located in the upper sector of every session viewer tab.
- **Dynamic configuration** — right-click a column field in the **Analysis Grid** viewer to display a context menu that enables Message Analyzer to automatically configure a view **Filter** that is coded with a corresponding field value.

  For example, you could right-click an address value in the **Analysis Grid** grid under the **Destination** column and then select the **Add "Destination" to Filter** command in the context menu that appears, to encode the right-clicked Destination address in a new Filter Expression, which might look similar to the following when complete: `WFPCapture.Destination == 192.168.1.1`, that is, if you used the **Microsoft-Pef-WFP-MessageProvider** to capture the messages.
- **Manual configuration** — manually configure a Filter Expression in the *filter expression* text box on the Filtering toolbar.

  To learn more about manually configuring Filter Expressions, see Writing Filter Expressions.

Using the Filtering Toolbar to Apply and Remove a View Filter

After you specify a Filter Expression on the Filtering toolbar for the data in a particular session viewer tab and you click the **Apply** button on the toolbar, the filter isolates the message data that meets the filtering criteria that you specified, such that only those messages are visible. However, a view **Filter** does not modify the original data set, as Message Analyzer persists the full data set such that you can redisplay it whenever you remove the specified view **Filter**. In addition, when you click the **Remove** button on the Filtering toolbar to undo the results of a view **Filter** that you applied to a set of messages, the Filter Expression text in the text box on the Filtering toolbar remains unchanged. This enables you to retain the Filter Expression text should you decide to reapply it, as described in Working with Tiered Filtering Configurations. For example, if you have a tiered configuration of two or more Filter Expressions, you can alternately select or unselect each filter to enable or disable it, respectively, and obtain different results based on different combinations of filtering criteria, for enhanced analysis.

> **NOTE**
>
> After you apply a view **Filter**, Message Analyzer indicates the number of messages that passed the filtering criteria by providing an indication next to the **Available** label in the lower-left sector of the Message Analyzer user interface.

Whenever you apply a built-in or manually-configured view **Filter**, Message Analyzer conveniently persists the Filter Expression text that you specified for future use. This enables you to access up to the last ten applied **Filters** by clicking the **History** drop-down list on the Filtering toolbar and selecting a chosen filter. In the **History** list, previously applied filters are displayed in chronological order, with the most recent applied filter as the first one in the list.

Assessing the Built-In Filter Expressions

Every Message Analyzer installation provides built-in Filter Expressions in a centralized **Library** that contains the **Message Analyzer Filters** asset collection, which is updateable from the Asset Manager dialog. In the **Library**, these Filter Expressions are contained in 12 different categories, in addition to an **Examples** category that contains an example filter that you can modify as required, so you can get started with developing Filter Expressions of your own. These categories and the filters they contain are described in the sections that follow. You are advised to assess these filters prior to applying them so that you will be able to more clearly recognize their effects upon trace results.

The categories in which the built-in Message Analyzer Filter Expressions are contained are described in the following sections:

- Azure Storage Category
- Address Filtering Category
- Diagnosis Category
- General Examples Category
- RegEx Category
- Contains Filters Category
- HTTP Category
- TCP Category
- LDAP Category
- Remove Noise Category
- File Sharing Category
- USB Category
- Examples Category

> **IMPORTANT**
> The **Azure Storage** category filters are not described in this section, as Microsoft provides related information in Azure blogs that you can access, as indicated immediately below.

More Information
To learn more about troubleshooting Azure storage logs, see the following documentation:
End-to-End Troubleshooting Using Azure Storage Metrics and Logging, AzCopy, and Message Analyzer
Monitor, Diagnose, and Troubleshoot Microsoft Azure Storage

Address Filtering Category

This filtering category contains several filters that isolate data based on Ethernet, IPv4, or IPv6 addresses, as follows:

- ***Address==02-01-0A-01-01-64** — filters out all messages except those that are intended for a specified MAC address, such as that of an Ethernet adapter.
- **IPv4.Address in 10.1.0.0/16** — filters for a particular IPv4.Address in a subnetwork generated by a specified subnet mask, which in the example is indicated by 16 (=255.255.0.0).
- **IPv4.Address == *192.168.1.1*** — filters out all messages except those that contain a **Source** or **Destination** address that matches the specified IPv4 address value.

> **NOTE**
> Before applying any of these filters, you should substitute your actual search address for the *italic* values in these Filter Expressions. The same is true of any built-in Filter Expression that has placeholder values.

- **IPv4.Address ~= *192.168.1.1*** — filters out all IPv4 messages that contain a **Source** or **Destination** address that matches the specified IPv4 address value. Use of the tilde (~) character ensures that you will only return IPv4 traffic. Otherwise, other traffic would not be blocked.

  For example, if you specified the Filter Expression as `IPv4.Address != 192.168.1.1` or `!(IPv4.Address == 192.168.1.1)`, all protocol messages that do not meet this criteria are returned, which means you would get all messages that are not IPv4 messages in addition to the target IPv4 traffic. However, as an alternative, you could add " `&& IPv4` " to either expression if you want to pass only IPv4 messages.

> **NOTE**
> The Filter Expression `!(IPv4.Address == 192.168.1.1)` is semantically equivalent to `IPv4.Address != 192.168.1.1`. For more details about the use of the "!" operator, see Other Filtering Considerations.

- **IPv6.Address == *2001:4898:0:FFF:200:5EFE:4135:4A7*** — filters out all messages except those that contain a **Source** or **Destination** address that matches the specified IPv6 address value.
- ***Address == *192.168.1.1* or *Address == *2001:4898:0:FFF:200:5EFE:4135:4A7*** — filters out all messages except those that contain a **Source** or **Destination** address that matches either the specified IPv4 address or IPv6 address value. *Address can also specify an Ethernet address.

> **NOTE**
> By using the expression "*Address", you can mix IPv4 and IPv6 addresses in this Filter Expression.

Diagnosis Category

This filtering category enables you to identify messages that have parsing errors. There are two types of filters in this category that reflect OPN parsing errors. These consist of **DiagnosisLevels**, which indicate parsing error severity level; and **DiagnosisTypes**, which describe the source of errors that can occur during parsing:

- **#DiagnosisLevels** — this filter passes all messages that have a diagnosis-level error. This Filter Expression is semantically equivalent to `#DiagnosisLevels != nothing`. If you want to be more specific, you can specify different enumeration values to filter for certain diagnosis levels. In the table that follows, friendly enumeration names are given along with their equivalent integer values that you can include as right-hand side values in this Filter Expression:

Table 15. Enum Values for DiagnosisLevel Filters

FILTER EXPRESSION WITH FRIENDLY ENUM NAME	FILTER EXPRESSION WITH INTEGER ENUM VALUE	DESCRIPTION
#DiagnosisLevels==Standard.DiagnosisLevel.Error	#DiagnosisLevels==1	Filters for messages that contain an Error-diagnosis level.
#DiagnosisLevels==Standard.DiagnosisLevel.Warning	#DiagnosisLevels==2	Filters for messages that contain a Warning-diagnosis level.
#DiagnosisLevels==Standard.DiagnosisLevel.Information	#DiagnosisLevels==4	Filters for messages that contain an Information-diagnosis level.

> **NOTE**
> In the Filtering Language, **DiagnosisLevels** is a global annotation, meaning that it is applicable to any OPN message. The character "#" provides access to an annotation, as differentiated from other filters that use a dot (.) operator to enable access to a field.

- **#DiagnosisTypes==2** — this filter passes any messages that contain validation type diagnosis errors that occurred during the OPN message ValidationCheck, as part of the parsing process. A validation error is a *soft* error, meaning that it is not severe enough to halt the parsing process. Validation errors typically occur when a message deviates from associated protocol specifications, which can include occurrences such as out-of-range values, invalid collection sizes, data element constraint violations, and so on.

  DiagnosisTypes is described in the OPN language as a flag pattern, which is a special type of enum. The table that follows specifies the enumeration values of the **DiagnosisTypes** enum that you can use in a **DiagnosisTypes** filter. Friendly enumeration names are given along with their equivalent integer values. You can specify the enum values in either format as right-hand side values in this Filter Expression.

Table 16. Enum Values for DiagnosisType Filters

FILTER EXPRESSION WITH FRIENDLY ENUM NAME	FILTER EXPRESSION WITH INTEGER ENUM VALUE	DESCRIPTION
#DiagnosisTypes == Standard.DiagnosisType.Application	#DiagnosisTypes==1	Filters for messages that contain Application-type diagnosis errors, for example, an application-related or network communication issue.
#DiagnosisTypes == Standard.DiagnosisType.Validation	#DiagnosisTypes==2	Filters for messages that contain Validation-type diagnosis errors. A Validation error is an indication that a message does not align with its protocol definition.
#DiagnosisTypes == Standard.DiagnosisType.InsufficientData	#DiagnosisTypes==4	Filters for messages that contain InsufficientData-type diagnosis errors. An InsufficientData error is an indication that message data was lost, for example, when Message Analyzer is attempting to group messages into an operation or when performing data reassembly.

FILTER EXPRESSION WITH FRIENDLY ENUM NAME	FILTER EXPRESSION WITH INTEGER ENUM VALUE	DESCRIPTION
#DiagnosisTypes == Standard.DiagnosisType.Parsing	#DiagnosisTypes==8	Filters for messages that contain Parsing-type diagnosis errors. A Parsing error is an indication that parsing failed when Message Analyzer attempted to decode invalid message data.

General Examples Category

This filtering category contains an example of how to AND two expressions together and an example of how to specify an explicit path to a field value in the TCP message hierarchy:

- **HTTP and !UDP** — this filter passes all messages that do not have HTTP over UDP.
- **TCP::Flags:SYN == true** — this filter uses an explicit path expression to pass TCP messages that have their SYN bit set (0x02). Note that the missing value between the double colons (::) is the **Segment** message-type specifier.

> **TIP**
> You could also specify the equivalent of this Filter Expression as: "`TCP:Segment::SYN==true`" or even "`TCP:::SYN==true`", to obtain an identical filtering result. At a minimum, an explicit path expression should specify a protocol or module as the first entity in the left-side expression and the field for which you are searching in a particular message hierarchy as the last entity in the expression, separated by an appropriate number of colons to delimit any skipped hierarchical entities. However, for the best performance, you should specify all entities in an explicit path.

RegEx

Message Analyzer provides several regular expressions (RegEx) that you can use as view **Filters**, as follows:

- **regex "(?!000)([0-6]\\d{2}|7([0-6]\\d|7[012]))([-]?)(?!00)\\d\\d\\3(?!0000)\\d{4}"** — enables you to find a phone number pattern.
- **regex @"^([\w\\.\\-]+)@([\w\\-]+)((\\.(\w){2,3})+)$"**— enables you to find an email address pattern.
- **regex @"^\d{3}-\d{2}-\d{4}$"** — enables you to find a social security number.
- **regex @"\bthis\W+(?:\w+\W+){1,6}?that\b"** — enables you to find the string "this" near the string "that".

Contains Filters Category

This filtering category contains filters that you can use to search for various strings and values in your trace results. You can search for messages that contain specified text in case-sensitive and case-insensitive searches. You can also locate hexadecimal patterns with this type of filter:

- ***Summary contains "error"** — this filter searches the **Summary** column for all messages in the **Analysis Grid** viewer that contain the text "error".
- **contains "Microsoft" caseSensitive** — this filter performs case-sensitive filtering that passes all messages containing the specified "Microsoft" search text. The default insensitive encoding will match ASCII, Unicode, or any other encoding type.

- **contains "Microsoft" encoding ASCII** — this filter performs case-sensitive filtering that passes all messages containing the specified search text in the supplied query encoding. The query encoding value can be any of the following: ASCII, UTF7, UTF8, Unicode, UTF32, BigEndianUnicode, and Base64.
- **contains "Microsoft"** — this filter performs case-insensitive filtering that passes all messages containing the specified search text.
- **contains $[*534d42*]** — this filter passes messages that contain the specified hexadecimal pattern, which in this example represents the ASCII characters for "SMB".

HTTP Category

This filtering category contains filters that isolate HTTP messages with errors, specified address text, and a specific stack configuration.

- **HTTP.StatusCode >= 400** — returns HTTP client and server error messages that have a StatusCode that is greater than or equal to 400, for troubleshooting purposes.
- **HTTP.Uri contains "msn"** — isolates HTTP addresses that contain the text "msn". The filtering criteria are both case- and encoding-insensitive.
- **HTTP\TCP.Port == IANA.Port.HTTP** — passes messages where HTTP is directly above TCP, as defined by the stack path specification "`HTTP\TCP.Port`". By using the stack path specifier "\", you can increase the specificity of your Filter Expressions.

TCP Category

This filtering category contains various Transport Layer filters that isolate TCP messages based on various criteria such as TCP ports, flags, window size, and options, as follows:

- **tcp.port == IANA.Port.HTTP** — enables you to filter for HTTP over TCP traffic in either direction. You can also substitute a port number in this expression, for example, "tcp.port == 80".
- **TCP.SourcePort in [*6608* , *6609*, *6610*]** — this filter passes all TCP messages that have a source port of 6608, 6609, or 6610. Note that the *italic* source port numbers in this expression are unassigned ports that serve as placeholders only.
- **tcp.options** — this filter passes only the TCP messages that have TCP options defined.
- **tcp.syn == true** — this filter passes TCP messages that have their SYN bit set (0x02). If you want to pass all TCP messages that specify the SYN field, you can drop the Boolean evaluation and apply the Filter Expression as: "`tcp.syn`".

> **NOTE**
>
> Although there is a **Segment** message type and a **Flags** container in the upper hierarchy for the **SYN** flag, it is unnecessary to specify a fully qualified expression such as "`TCP.Segment.Flags.SYN`". In the Filtering Language, the dot notation is commonly used to traverse the message type-hierarchy, which is shown in the **Field Chooser Tool Window**. A dot (.) in a Filter Expression means the filter should find any entity at any depth in the hierarchy, hence in this example: `tcp.syn==true`. However, to improve performance, you can create an explicit path that uses colons, as indicated in the text of the earlier Filter Expression `TCP::Flags:SYN == true`.

- **TCP.Windowscaled < *1000*** — this filter passes TCP messages having a receive window size that is less than 1000.

  Note that TCP Window Scaling increases the TCP receive window size above its default maximum value of 65,535 bytes for better throughput, for example, in cases where a receiver is overwhelmed with high

message volumes such as a web server might be.

> **NOTE**
>
> If you encounter a bottleneck, it might be caused by a small TCP receive window size that should be increased above the default setting by using a tool such as **netsh**. Windows 7, Windows 8, Windows 8.1, and Windows Server 2012 operating systems provide a Receive Window Auto-Tuning Level configuration that is set by default to "Normal", which allows the TCP receive window size to grow automatically to accommodate most conditions. If TCP Auto-Tuning is disabled, the receive Window size is limited to the default size of 65,535 bytes (64 KB).

LDAP Category

This filtering category contains an LDAP filter that removes all traffic that transits the LDAP port.

- ***Port != IANA.Port.LDAP** — removes all LDAP port traffic, regardless of the protocol. Port represents either SourcePort or DestinationPort for UDP and TCP messages, or any protocol with a field or property named Port. The asterisk in the expression "*Port" indicates that the filter should look at any protocol that defines a field or property named "Port" and filter out any LDAP port traffic that might have been invoked by that protocol. This filter will also pass messages from any protocol where a Port field is nonexistent.

Remove Noise Category

This filtering category contains filters that remove unwanted messages so that you can display a more streamlined message set.

- **!(*Port in [3389, 1494, 1503])** — removes all Remote Desktop messages that traverse ports 3389, 1494, and 1503 in either direction, for various remote desktop (RDP) communication protocols.
- **WiFi.FrameControl.Type !=0** — removes Wi-Fi noise by filtering out beacons.

File Sharing Category

This filtering category contains several examples that you can use to filter trace results for SMB messages with diagnosis errors, SMB messages with a **Filename** field defined, or any file sharing traffic captured on TCP port 445:

- ***SMB.FileName ~= "" OR SMB2.FileName ~= ""** — this filter passes only SMB and SMB2 messages that have a FileName field defined, in which the field value is not equal to an empty string. This Filter Expression is not the same as `SMB.FileName != "" OR SMB2.FileName != ""`, which returns non-SMB traffic as well. By using the tilde (~) character in this filter, you ensure that only SMB traffic is returned.
- **SMB#DiagnosisLevels** — this filter passes all SMB messages that have parsing errors. This filter is equivalent to the expression `SMB && #DiagnosisLevels`. You can also apply any of the enumeration values specified by the table in Diagnosis Category to this Filter Expression. For example, you could specify `SMB#DiagnosisLevels== Standard.DiagnosisLevel.Warning` to return all SMB messages that contain a Warning diagnosis level.
- ***Port == IANA.Port.SMB** — this filter passes messages from any protocol that has a top-level source or destination Port field equal to 445, including TCP and UDP. It also returns any top-level protocol that has a field or property named "Port".

USB Category

This category contains an example of filtering for USB transfer messages that contain errors:

- **Microsoft_Windows_USB_USBPORT.fid_URB_Hdr_Status ~= 0** — this filter passes all messages from the **Microsoft-Windows-USB-USBPORT** provider that contain the Fid_URB_Hdr_Status field with a value not equal to zero, indicating an error condition. Note that this Filter Expression is not the same as `Microsoft_Windows_USB_USBPORT.fid_URB_Hdr_Status != 0`, which returns all messages captured by the provider, whether it had the specified field or not.

Examples Category

The **Examples** filtering category exists under **My Items**, where all view **Filters** that you create and save with the **Edit Filter** dialog are displayed. The **Edit Filter** dialog is accessible by clicking the **New Filter** item in the **Library** drop-down list in any Filter panel on the Filtering toolbar. The **Examples** category contains a single filter that you can modify as you like, so that you can get started with creating your own filters.

> **NOTE**
>
> Any Filter Expression that you create becomes part of your centralized user **Library** and is available to you as **Session Filter** whenever you are configuring a Data Retrieval Session or Live Trace Session. Such a Filter Expression is also available to you as a view **Filter** whenever you are analyzing trace results in an Analysis Session. In addition, note that Filter Expressions you create can contain an **Alias** or a **Union**, as appropriate.

The Filter Expression that is included as a startup example is specified as follows:

- **HTTP** — this filter will return all top-level HTTP operations and any other messages where HTTP exists in the stack.

  You are encouraged to browse through the HTTP message hierarchy with the **Field Chooser** window to review the types of fields that you can specify in an HTTP filter. You can also make use of the Filtering IntelliSense feature that kicks in and exposes the message hierarchy when you type a dot (.) after "HTTP" in the Filter Expression text box of the **Edit Filter** dialog (or in the Filter Expression text box on any Filter panel). You can open the **Edit Filter** dialog by right-clicking the HTTP example and selecting the **Edit** item in the context menu that appears.

Sharing Filter Items

You can share items from your centralized filter **Library** with other users, including any that you create, by exporting them as a **Filter** asset collection (*.asset file) through the Message Analyzer Sharing Infrastructure. The features of the Sharing Infrastructure enable you to export and import **Filter** items to and from a designated user file share or other location, respectively, or to a user feed that you create from the **Settings** tab of the Message Analyzer **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu. To learn how to manage view **Filters** from the **Manage Filters** dialog, see Applying and Managing Filters.

More Information

To learn more about using **Aliases** in Filter Expressions, see Performing Message Analyzer Operations with Aliases.

To learn more about using **Unions** in Filter Expressions, see Performing Message Analyzer Operations with Unions.

See Also

Writing Filter Expressions

Filter IntelliSense Service

Filtering Column Data

2 minutes to read

To enable you to rapidly isolate messages containing specific field values, names, or other text in a set of trace results displaying in the default **Analysis Grid** viewer, Message Analyzer provides a **Column Filter** feature. This feature enables you to filter a set of trace results based on search text that you enter in a **Column Filter** text box above most **Analysis Grid** viewer data columns. As soon as you enter a text value in a **Column Filter** text box, only the messages that meet the criteria of the search text that you specify in a particular column are displayed, while all other messages are temporarily hidden. To remove the effects of an applied **Column Filter** value and return to the original message set, you can manually clear the specified text value or you can click the **Clear Search** "x" in the **Column Filter** text box.

Enabling Column Filter Text Entry

To display **Column Filter** text boxes for the data columns displayed in the **Analysis Grid** viewer, click the **Show Column Filter Row** icon in the upper-left sector of the **Analysis Grid** viewer tab, immediately to the left of the **MessageNumber** column. When you click the **Column Filter** icon, a row of light amber-colored **Column Filter** text boxes displays immediately below the name of each column, with the exception of the **Diagnosis** column. To hide the **Column Filter** text boxes, click the **Show Column Filter Row** icon again.

> **NOTE**
>
> The **Column Filter** feature is also available in the message **Details Tool Window**. You can display these **Column Filter** text boxes in **Details** in a manner that is similar to the previously described method.

Locating Messages with Matching Search Text

As you type search text into a **Column Filter** box, the filter results display immediately so you can quickly locate messages, field values, or other text of interest. For example, if you ran a Live Trace Session with the **Local Network Interfaces Trace Scenario** and you wanted to look at all the UDP messages that you captured, you could type "UDP" in the **Column Filter** text box at the top of the **Module** column of the **Analysis Grid** viewer. This action would display only the top-level UDP messages in a trace. Note that **Column Filters** do not recognize text matches in the origins tree (stack), unless a stack message node containing the search text is already open and the matching text is exposed. In this case, such a message would also be returned as a match.

> **TIP**
>
> **Column Filter** input values are case-insensitive.

Writing Filter Expressions

2 minutes to read

This section provides you with practical knowledge that you can utilize to create your own Filter Expressions with the Message Analyzer Filtering Language. This language is a derivative of the full OPN language that developers use to write OPN protocol descriptions that Message Analyzer utilizes to parse the messages of such protocols. As a result, there are various elements, constructs, and syntax that both languages share. Any Filter Expressions that you create can be saved in the centralized filter **Library**. Thereafter, you can use them as a **Session Filter** when configuring a new session, or as a view **Filter** when analyzing trace results in an Analysis Session.

To aid your understanding of the Filtering Language, you are advised to work with some of the built-in Filter Expressions in the centralized filter **Library** in parallel with the studies that you do here. The goal is to accelerate the learning process through observation and assessment of working example results. You can even work with any existing Filter Expression by creating a copy, modifying the copy, and saving it as a new Filter Expression in your central user **Library** without affecting the original filter configuration. This enables you to start with a known working filter configuration and perform some experimental modifications to see the effects of your changes when applied to a set of trace results.

Although you can specify a Filter Expression as a **Session Filter** when configuring a Data Retrieval Session or a Live Trace Session, the topics in this section focus on applying view **Filters** in an Analysis Session with the default **Analysis Grid** viewer, because this environment provides the most robust environment for demonstrating filtering functionality.

What You Will Learn

In this section, you will learn about the Filtering Language and how to create some simple Filter Expressions, in addition to others that are more complex. The material begins with an introductory overview to help you get started with custom Filter Expressions and includes a few simple examples followed by a description of the Filter IntelliSense Service. The discussions then move into understanding the basics about the Filtering Language, which includes the use of logical operators, arithmetic operators, and literals; how to traverse message hierarchies; and applying other filtering considerations. A walkthrough of filter features, special functions, and other capabilities concludes this section. In addition, examples are provided that combine different language features, including various applications of Filter Expression syntax, semantics, statements, truncation, traversers, aliases, and so on. This content is covered in the topics below:

Introduction to Creating and Applying Filters
Understanding the Filtering Language Basics
Using the Filtering Language

See Also

-Procedures: Using the Data Filtering Features

Introduction to Creating and Applying Filters

5 minutes to read

If the built-in Filter Expressions in the centralized filter **Library** do not provide the filtering functionality you need to analyze your message traffic, you can create your own custom Filter Expressions. To help you create Filter Expressions, Message Analyzer has a built-in Filter IntelliSense service that provides a statement completion capability that is similar to the Visual Studio Statement Completion Service. The Filter IntelliSense Service provides support for configuring **Session Filters** for a Data Retrieval Session or a Live Trace Session, and for configuring view **Filters** in an Analysis Session. For further details about Filter IntelliSense for Filter Expressions, see the Filter IntelliSense Service topic.

Creating a Custom Filter

To create your own custom **Filters**, you will need to understand the Message Analyzer Filtering Language. As indicated earlier in Assessing the Built-In Filter Expressions, you should first familiarize yourself with various Filter Expressions that exist in the **Message Analyzer Filters** asset collection, which is accessible from the **Library** drop-down list on the Filtering toolbar that appears above the main analysis surface for any set of trace results. As you work with these filters and observe how they impact your trace results, you can derive an understanding that will help you when creating your own Filter Expressions that you can use as a view **Filter**, **Session Filter**, or **Viewpoint Filter**, which all use the same centralized filter **Library**. In the sections listed immediately below, you will learn about the concepts and constructs of the Filtering Language, so you can use this information to create custom Filter Expressions and optionally add them as new assets to the centralized filter **Library**.

Generating a View Filter Dynamically

The easiest way to create a Filter Expression and apply it to message data displaying in the **Analysis Grid** viewer is to right-click a **Column** value and select the **Add <*columnName*> to Filter** command, to add the **Column** value as a Filter Expression. For example, if you right-click a protocol module in the **Module** column, the protocol name is added to the Filter Expression text box in a Filter panel on the Filtering toolbar as your Filter Expression. This action creates a simple filter that is also referred to as an "atomic" filter, meaning that the Filter Expression contains no logical combinators such as OR, AND, or NOT, and has no left-hand-side expression. Note that dynamic generation of a Filter Expression only creates the filter code; you must then apply the Filter Expression to see the effects, as described in the section that follows.

Selecting and Applying a View Filter

As described in many topics of this Operating Guide, you can *select* a built-in Filter Expression from the centralized **Library** to apply to a set of trace results, as appropriate to the type of analysis you are performing. When you are ready to *apply* a view **Filter** to your trace results, click the **Apply** button on the Filtering toolbar to see the effect of the **Filter** on your message data. If you added a protocol name as your Filter Expression as previously indicated, the **Analysis Grid** viewer will display messages from the specified protocol only, which includes top-level messages or operations with that protocol name, in addition to any other top-level message that contains that protocol in its origins tree (stack).

Note that this differs from the way **Column Filters** work in regard to stack messages, where filtering evaluates exposed layers only, as described in Filtering Column Data. The application of a view **Filter** is extremely useful for rapidly isolating the message data you want to analyze, especially in very large traces with thousands of messages. Also, in using the right-click method for adding a **Filter**, configuration is fast and return results are guaranteed, given that these Filter Expressions are derived from values that already exist in the set of trace results with which

you are working.

> **NOTE**
> Whenever you apply a Filter Expression from the Filtering toolbar, the in-focus session viewer tab and the corresponding session viewer node in **Session Explorer** both provide an indication of the Filter Expression that has been applied. You can view these indicators by hovering over the appropriate tab or node with your mouse. You will also see a funnel icon in both of these locations to indicate that a Filter Expression has been applied.

You can also use a Filter Expression to restrict your filter results to messages containing specific values. For example, you might specify a Filter Expression such as `UDP.Length > 100` to return only those UDP messages that have a message Length (Header + Payload) value that is greater than 100 bytes. Also, if you want to display all messages that are not UDP traffic, you could apply a filter such as `!UDP`. Moreover, if you want to filter out all UDP and TCP messages, you could apply a filter expression such as `!UDP&&!TCP`.

> **IMPORTANT**
> The scope of a view **Filter** is NOT global, meaning that it's effects are applied to only the data viewer that is currently in focus. All other data viewers that are out of focus are not impacted. Therefore, you must apply separate **Filters** to each data viewer independently, as appropriate.

Maintaining History of View Filter Application

After you create and apply a custom Filter Expression, it is maintained in a common **History** drop-down list, where the last filter applied appears first. Like the centralized filter **Library**, the **History** drop-down list is accessible from any Data Retrieval Session, Live Trace Session, or Analysis Session, so that you can reapply any custom Filter Expression in the common **History** drop-down list as a **Session Filter** or a view **Filter**.

> **NOTE**
> The **History** drop-down list currently maintains the last set of ten filters that you applied. However, for more permanent storage, you can save any custom Filter Expression that you create to the centralized filter **Library**, if you find that it is useful. In addition, you can share all **Filter** items that exist in the **Library** with others through the Message Analyzer Sharing Infrastructure.

In the remaining topics of this section, you will learn more about the Filter IntelliSense service in addition to the concepts, semantics, syntax, and other details about the Filtering Language, so you can create robust Filter Expressions that manipulate trace results in unique ways for data analysis purposes. You can also review a set of procedures that provide numerous examples of using Filter Expressions to isolate data for focused analysis, as indicated below in ***See Also***.

See Also

Procedures: Using the Data Filtering Features
Applying and Managing Filters

Filter IntelliSense Service

6 minutes to read

To assist you in developing Filter Expressions, Message Analyzer provides a statement completion service known as Filter IntelliSense. The Filter IntelliSense service is a level-sensitive and interactive feature that interprets the text that you enter in Message Analyzer Filter Expression text boxes and responds by displaying scrollable drop-down lists of message elements, such as protocols, message types, structures, fields, properties, annotations, and so on. This feature helps you to navigate through message hierarchy levels to find and select elements that you can configure in Filter Expressions. Similar to the **Field Chooser Tool Window**, the Filter IntelliSense service enables you to quickly learn about and work with the message hierarchies of the protocols that are parsed by the PEF Runtime. Because Filter IntelliSense provides visual message element cues in response to text that you enter, it enables you to discover and learn about the Filtering Language syntax on-the-fly as you use Message Analyzer to resolve issues on which you are working. The visual presentations provided by the service also help you to easily recognize the difference between ambiguous and fully qualified filter expressions, such as `HTTP.Method=="GET"` and `HTTP.Request.Method=="GET"`, respectively.

Starting the Filter IntelliSense Service

There are several ways that you can initiate the Filter IntelliSense service. First, you can type specific characters in a blank Filter Expression text box to start the service. For example, you might type the name of a protocol in the text box of any Filter panel on the Filtering toolbar during session results analysis. You can also start the service by initiating the Keyboard shortcut **Ctrl+Spacebar** while the cursor is located in any blank Filter Expression text box.

The following list describes the results of starting the Filter IntelliSense service with various actions that include character entry in a Filter Expression text box and the keyboard shortcut method:

- **Enter a valid OPN-qualified identifier character** — displays a list of elements in alphabetical order that match the input text character(s).
- **Enter the "#" character** — displays a list of global annotations in alphabetical order.
- **Enter the "*" character** — displays a list of filter expressions that can use the wildcard character, in alphabetical order.
- **Press Ctrl+Spacebar** — displays a list of protocols and global annotations in separate groups, with each group in alphabetical order.

> **NOTE**
>
> To place the cursor in any Filter Expression text box, simply click directly in the text box of choice. This includes the text box in any Filter panel on the Filtering toolbar during session analysis, or the **Session Filter** text box that exists in the **New Session** dialog during session configuration.

Using the Filter IntelliSense Service

When you invoke Filter IntelliSense in one of the indicated ways, each listed entity displays with an associated icon that represents its element type, for example, a protocol module, message, structure, operation, property, or field, similar to the way the **Field Chooser** window displays icons to the left of each element. If you select any listed element, the element type, its data type, and one or more fully qualified expressions are displayed in an informational pop-up window to the right of each element.

If you continue to type with Filter IntelliSense invoked, the element list is filtered to show only the results that match the input text. If you invoke Filter IntelliSense in a blank Filter Expression text box, the results are shown in separate sections for each element type with each section in alphabetical order. If you invoke Filter IntelliSense from below a parent element, it displays child entities and its descendants in separate groups with each group in alphabetical order. If a Filter Expression is abbreviated or ambiguous and can be represented by two or more expressions, a list of represented, fully-qualified expressions will display in an informational pop-up window.

For example, when you type one or more characters in any Filter Expression text box, Filter IntelliSense tracks the text entries that you enter and responds with a list of elements that match the typed characters, for example, matching protocol or field names. Filter IntelliSense also highlights an element in the list that matches the input character to enable you to conveniently add that selected element to your Filter Expression by pressing **Tab**, **Enter**, or the **Space bar** on your Keyboard. If the selected element is not appropriate for your Filter Expression, you can scroll down the list to find another one by using the down arrow key.

If you invoke Filter IntelliSense at root level by typing a protocol name followed by a dot (.), a list displays that can include message types, structures, properties, fields, and so on, depending on the protocol and its message hierarchy. From this level, you can also select a field or other message element from the list and integrate it into your Filter Expression by pressing one of the previously specified Keyboard keys. If you type the name of a field only, a pop-up list displays with all possible filter expressions that end with the name of the field you specified. As you continue to traverse the message hierarchy of a particular protocol by using the dotted notation, Filter IntelliSense continues to display the field names that are associated with the current hierarchic level. This process repeats until you reach the last available level in the protocol message hierarchy.

Service Implementation

To implement its functions, the Filter IntelliSense service makes use of a metadata access layer that is made available by underlying classes that provide the routines for fetching statement completion text. The service architecture that underlies these functions includes the following components:

- Statement Completion Service with a cache
- Statement Completion Store
- Query Name Reference Resolver

The presentation of statement completion text in Message Analyzer Filter Expression text boxes occurs in response to the initiation actions specified in Starting the Filter IntelliSense Service. Note that the Filter IntelliSense service is available for view **Filter** configuration on the Filtering toolbar, even while Message Analyzer is still parsing messages.

Interactive Intelligence
The Filter IntelliSense service contains the interpretive intelligence to anticipate the elements for which you might be searching based on the characters you enter. For example, if you enter the text "Port" in a Filter Expression text box, the service matches the prefix and entities such as "PortName" display in the results list; however, element names such as "DestinationPort" do not display. In addition, the portion of the prefix text that matches the entered text is displayed in bold.

If you invoke Filter IntelliSense at the root level in an empty Filter Expression text box, the results are grouped separately and listed in the following order, with each group alphabetically sorted:

- Protocol modules
- Global annotations

If you invoke Filter IntelliSense at the root level in an empty Filter Expression text box and *then* enter text to activate the list filtering feature, the results are grouped separately and listed in the following order, with each

group alphabetically sorted:

- Protocols
- Global annotations
- *Messages, *Structures, *Fields, *Properties
- Messages, Structures, Fields, Properties

If you invoke Filter IntelliSense below a parent element at root level, for example by entering the text "TCP." in a Filter Expression text box, the following child entities are listed together in a single group that is alphabetically sorted:

- Messages
- Structures
- Fields
- Properties

If you continue to invoke Filter IntelliSense at the child level, lists of fields display that are associated with the various child entities. From the lists, you can select the message elements you require for integration into your Filter Expression.

> **NOTE**
> The Filter IntelliSense service does not respond to colon (:) characters that you might normally enter to create explicit paths in Filter Expressions. It also does not respond to other characters such as operators and backslashes, the latter of which you can use to traverse a protocol stack.

More Information
To learn more about explicit paths, see Traversing the Message Hierarchy with Explicit Paths.
To learn more about special methods for traversing a protocol stack, see Browsing Message Origins.
To learn more about annotations, see Accessing Message Properties and Annotations.

See Also

Using the Filtering Language

Understanding the Filtering Language Basics

8 minutes to read

To assist you in understanding the Filtering Language, you should experiment with some of the built-in Filter Expressions that Message Analyzer provides in the centralized filter **Library**. You should also review the descriptions of these filters in the Filtering Live Trace Session Results topic of this documentation and then apply them to a trace, to get a sense of how real-world filter expressions work when analyzing trace results.

This section provides concepts and constructs that you will need to understand to create your own Filter Expressions. In this section, you will learn how to use the basics of the Filtering Language to construct Filter Expressions. This includes the use of operators and literals, traversing the protocol message hierarchy, and other considerations such as case sensitivity, filter applicability, semantic equivalence, and the meaning of negation.

Using Operators

When creating custom Filter Expressions, you can use the basic Boolean operators that follow. You can also use the textual equivalent of these operators, for example, AND, OR, and NOT in your filter expressions:

- **&&** — represents the logical AND function. Typically used when combining filter expressions.
- **||** — represents the logical OR function. Typically used when combining filter expressions.
- **!** — represents the logical NOT function. Typically used for negation.

You can also apply the following relational operators to applicable Filter Expressions when you need to restrict fields to specific values:

- **==** — Equals. An operator that evaluates two filter expression operands for value equality.
- **!=** — Not equals. An operator that evaluates two filter expression operands for value inequality. Note that this operator also evaluates nonexistence as a form of negation.
- **~=** — Not equals. This operator negates the condition on a value only but does not evaluate nonexistence as a form of negation.

> **NOTE**
>
> The following examples further clarify the difference between the "!=" and "~=" operators. A filter expression such as `TCP.SourcePort != 443` returns all TCP messages that have a **SourcePort** value that is not equal to 443, together with all messages that are not TCP. By using the ~= operator in this expression, for example `TCP.SourcePort ~= 443`, the condition on the value is negated and the filter expression will return TCP messages with a **SourcePort** that is not equal to 443, but non-TCP messages will not also be included.

- **>** — Greater than. This operator is used to evaluate whether one filter expression operand is greater than the other.
- **>=** — Greater than or equal to. This operator is used to evaluate whether one filter expression operand is greater than or equal to the other.
- **<** — Less than. This operator is used to evaluate whether one filter expression operand is less than the other.
- **<=** — Less than or equal to. This operator is used to evaluate whether one filter expression operand is less than or equal to the other.

- **in** — Membership in an array, set, or map. This operator is used in a filter expression to determine whether a particular left-side field or literal value exists in a user-specified right-side collection of values, such as an array. For example: `IPv4.Address in [192.0.1.1, 192.0.0.0, 192.0.0.2]` or `TCP.SourcePort in [6608, 6609, 6610]`.

More Information
To learn more about how to apply relational operators, see the built-in Filter Expressions described in Filtering Live Trace Session Results.

The Filtering Language also supports the following bitwise and arithmetic operators that you can apply to field expressions:

- **Boolean and bitwise operators**: |, ^, &, ~, and !.
- **Bit shift operators**: << and >>.
- **Arithmetic operators**: +, -, *, /, and %.

For example, the following is a valid filter expression that uses arithmetic operators: `TCP.SourcePort + 1 == TCP.DestinationPort / 2`.

Using Literals

The Filtering Language supports all literals for built-in Open Protocol Notation (OPN) types, such as integer, floating point, Boolean, char, string, and so on. The language also supports protocol-specific literals, as indicated in the examples of the table that follows:

Table 17. Using Literals in Filter Expressions

LITERAL TYPE	EXAMPLE
IPv4	"127.0.0.1"
IPv4 subnet	"10.10.249.1/21"
IPv6	"2001:0db8:85a3:0000:0000:8a2e:0370:7334", "2001::A001"
IPv6 subnet	"2001: b8::/32"
MAC	"48-2C-6A-1E-59-3D"
Binary	"$[FFFE]"
GUID	"{448ea956-38ff-4620-b022-dd88cac6c896}"
Arrays	"[80, 339, 993]"

NOTE
When a field is declared as optional in OPN, the special value "nothing" is used to represent when the field is not present in an incoming message. You can include this special value in a filtering expression to represent such a case. For additional details about optional declarations and the use of the "nothing" value, see the OPN Programming Guide on the Microsoft downloads site.

Traversing the Message Hierarchy

The Filtering Language enables you filter messages based on different entities, such as protocol names, message types, structures, properties, or other field names. You can also filter for annotations, which generally fall into the "field name" category. As indicated earlier in this topic, you can also restrict the messages returned by an applied filter to specific values contained in their fields, and you can combine atomic filters with standard Boolean and relational operators.

If you want to create a filter that passes messages from a specific protocol that contains a particular field name or value, you can use the dot (.) notation to traverse the message type hierarchy by specifying a top-level protocol name, followed by the message type, and ending with a field. Whenever you need to traverse the message type and field hierarchy of a particular protocol using dotted notation, you can use the following fully-qualified expression pattern:

```
ProtocolName.MessageType.Field1.Field2. ...Fieldn
```

For example, in the following figure you see the hierarchy for TCP, which you can view in an Analysis Session when you display the **Field Chooser Tool Window**:

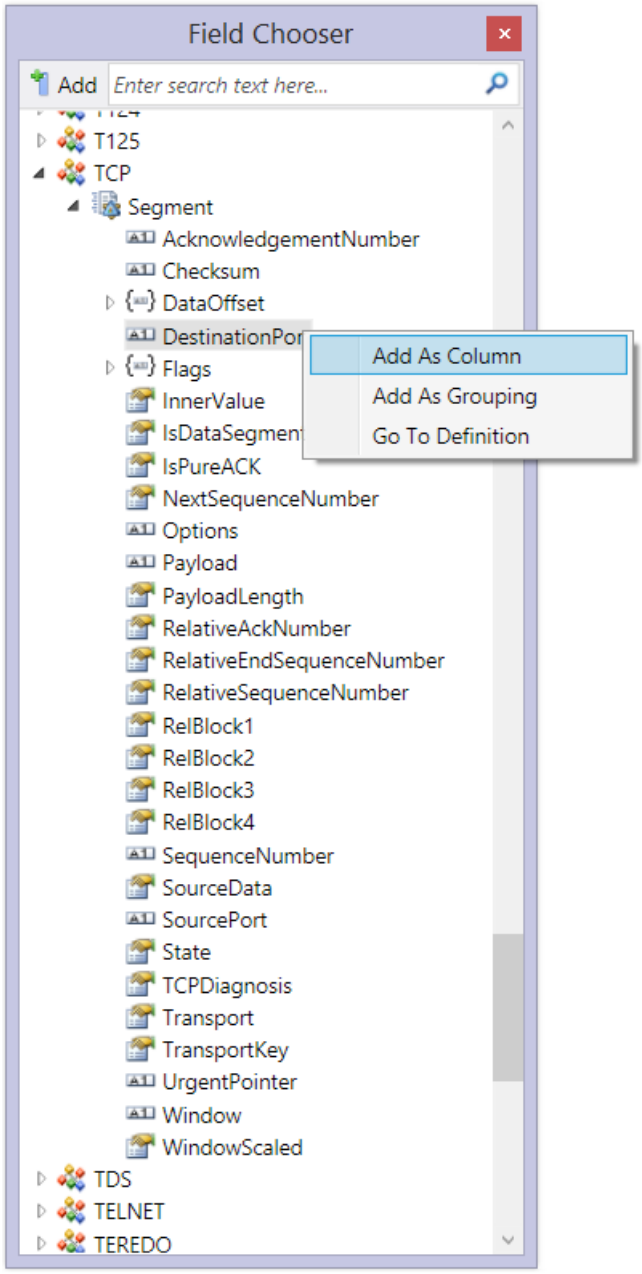

Figure 67: Field Chooser Tool Window – TCP message hierarchy display

With the TCP message hiearchy as a reference, you could inspect the hierarchy to view the available fields so that

you can write filters similar to the following:

- `TCP.Segment.SourcePort==443` — filters out all TCP messages except those that traverse port 443 (HTTPS).
- `TCP.Segment.Flags.Syn` — enables you to look at the state of the **SYN** bit for all TCP messages in your trace results.

  > **NOTE**
  > To examine these values, you can add a **SYN** column to your **Analysis Grid** viewer column configuration with the **Field Chooser**.

- `TCP.Segment.Flags.Syn==True` — filters out all TCP messages except those that have their **SYN** flag bit set to True (indicating that a TCP message is part of a TCP three-way handshake for connection setup).

> **NOTE**
> When you enter text in a Filter Expression text box to configure a Filter Expression, the Filter IntelliSense Service is activated. This service enables you to navigate message hierarchies on-the-fly, to locate specific message elements that you can integrate into your Filter Expression, as described in Filter IntelliSense Service.

> **TIP**
> Based on the highlighted field in the TCP message hierarchy shown in the previous figure, you can add the **DestinationPort** field as a new column to the **Analysis Grid** viewer (**Add as Column** command in the context menu), and as a new **Group** to the **Grouping** viewer (**Add as Grouping** command). To do this, the respective viewer will need to be in focus. You can also display the **DestinationPort** definition from the TCP.opn file in a separate viewing tab by choosing the **Go To Definition** command in the right-click **Analysis Grid** viewer context menu.

Other Filtering Considerations

When creating custom Filter Expressions, you should also consider the following:

- **Case sensitivity** — in general, filters are case-insensitive, unless case-sensitivity options are accepted for a particular expression type. For example, "UDP" and "udp" are functionally equivalent.
- **Filter applicability** — in the **Analysis Grid** viewer, messages are grouped as expandable, top-level parent nodes containing expandable, lower-level child-node messages that comprise the origin tree, as described in Viewing Message Data. These expansion nodes are each labeled by a **Message Number**. The child messages usually consist of the message protocol stack, the base ETW module, and can also include reassemblies such as you have with TCP virtual messages. When you apply a Filter Expression to a message collection, all messages that match the filtering criteria are returned, including all top-level messages that contain one or more origins tree messages where a match is found, regardless of whether the top-level messages specifically met the filtering criteria.

  For example, if you apply a filter expression such as "Ethernet", the filter will return all Ethernet messages in addition to upper-layer protocol messages such as IP, TCP or UDP, and so on, because these messages form part of the origins tree.

- **Semantic equivalence** — the semantics of the "!=" operator is equivalent to negating a filter that contains the operator "==". For example, for any expressions of A and B, the filter "!(A == B)" is entirely equivalent to "A != B". Therefore if you use the operator "!=", its implied negation produces the previously described complement set of messages. This means that a filter such as `!(TCP.SourcePort == 443)` is semantically equivalent to the filter `TCP.SourcePort != 443`.

- **Negation** — if you apply a hypothetical filter named "TCPFilter" to a message collection, the filter "!TCPFilter" would return the complement set of messages that "TCPFilter" returned from the collection. For a more specific example, if you apply a filter expression such as `TCP.SourcePort == 443`, it will return all TCP message traffic that traversed TCP port 443. If you negate and apply this expression as `!(TCP.SourcePort == 443)`, it will specifically filter out all TCP messages that have a **SourcePort** field with a value of 443. In addition, other messages that meet the negated filter criteria will be passed by this filter, including other protocol messages in the collection that do not have a TCP **SourcePort** field defined, which in this case would mean all messages that are not TCP. This is the case because negation is inclusive of nonexistence.

- **Precedence** — most operators that you use in a Filter Expression are designated with a predefined evaluation order of precedence. Generally, the use of parentheses in a Filter Expression helps to define how the expression is evaluated. However, without the use of parentheses, some ambiguity can be introduced in the interpretation of your Filter Expressions. Therefore, when you use the ! (NOT) and == (EQUALITY) operators in a Filter Expression, the compiler gives precedence to the == operator when evaluating the expression. For example, this would make the expression `!F==true` semantically equivalent to `!(F==true)`.

More Information

To learn more about OPN concepts that may help you write Filter Expressions, see the OPN Programming Guide document.

Using the Filtering Language

10 minutes to read

The topics of this section as listed below describe how to use various features of the Filtering Language.

Using Truncated Filter Expressions
Traversing the Message Hierarchy with Explicit Paths
Accessing Message Properties and Annotations
Referencing Enumerations
Using Special Filtering Functions
Browsing Message Origins
Using Aliases

For each of these discussion areas, working filter examples are provided so that you can copy and paste them into a Filter Expression text box and test them. However, please observe the following advisory note.

Caution

When copying and pasting filter expressions into a Filter Expression text box, please be aware that quotation marks may not work as expected due to environment formatting issues. If this occurs, simply retype the quotes in the Filter Expression text box and then apply the filter. Also note that other formatting issues may occur when copying content from HTML.

Using Truncated Filter Expressions

When writing a Filter Expression, you are not always required to provide the fully qualified expression to a field, since the dot notation is interpreted by the OPN Compiler to mean "look for a matching entity at any hierarchy depth". For example, the filter expression `TCP.Segment.Flags.Syn==True` can also be written `TCP.Syn==True` to return an identical result. In practice this means that if there is more than one possible match with such a filter, it will return all matches, which can introduce a measure of ambiguity. In addition, the Filtering Language allows you to use a placeholder asterisk to shorten the expression to `*Syn==True`. This filter expression would then look for *any* protocol that has a SYN field set to True, which for example could expectedly include a TCP message, but also could include a message such as a Session Multiplex Protocol (SMP) connection request, which also uses a SYN flag when making a session connection.

> **NOTE**
> If you do not specify a leading asterisk in your Filter Expression, a protocol name is required.

The following table depicts some of these variations in Filter Expressions, using the HTTP protocol as an example.

Table 18. Truncating Filter Expressions

EXPRESSION	MEANING
`HTTP.Method==“GET”`	This filter returns only HTTP messages that have the value of their Method field set to "GET". In this filter, note that the HTTP "Request" message type is not specified.

EXPRESSION	MEANING
`*Method==“GET”`	This filter has a result that is identical to `HTTP.Method==”GET”`, although both the protocol name and message type are omitted from the expression. Although this filter expression looks for messages from *any* protocol with a Method field set to “GET”, it is highly unlikely that such a message would be found, thus only HTTP messages will be returned.
`*Port==443`	This filter looks at any protocol that has Port field defined. TCP defines SourcePort and DestinationPort fields, so this filter will capture HTTPS over TCP traffic because “Port” is an alias that represents both these fields, as described in Using Aliases. However, messages from any other protocol that have a Port field defined might also be returned if their Port values are set to 443. **Note:** When viewing the results of a filter such as this, consider that if the origin tree includes messages with a Port value equal to 443, you may also see other messages with Port values that are not set to 443, as indicated in the **Filter applicability** bullet item in Other Filtering Considerations.

Traversing the Message Hierarchy with Explicit Paths

When you apply a Filter Expression that uses an asterisk, as in the example `*Port==443`, you might notice that additional processing time is incurred to return filter results and that there may be a level of ambiguity inherent to those results. To improve the efficiency of such a filter and to disambiguate the results, you can specify an explicit traverse path for the filter by substituting the colon (:) separator for the dot (.) notation. When you use a colon separator to traverse the message hierarchy, the OPN Compiler interprets the separator to mean “look for a matching entity exactly one level below”. The following table provides some examples of what several filter expressions might look like with the colon (:) separator.

Table 19. Using Explicit Traversal Paths in Filter Expressions

EXPRESSION	MEANING
`HTTP:Request:Method == “GET”`	This filter returns only HTTP Request messages that have their Method field set to “GET”. By using colon separators in this filter expression, any other entity in the HTTP message hierarchy that is named “Request” will be ignored.
`HTTP::Method == “GET”`	This filter also returns only HTTP Request messages that have their Method field set to “GET”. The missing value between the double colons is the “Request” message type specifier. The use of two separators signifies to the compiler that it should look for a matching entity two levels below.

Accessing Message Properties and Annotations

Other data that you might want to return with a Filter Expression includes message properties and annotations. You can access message properties the same way that you access message fields, by using the dot notation to traverse the message hierarchy. However, the Filtering Language allows you to access message annotations, by using the “#” operator. In the Filtering Language, an annotation is additional information that is not directly related to a message, such as user-provided comments, implementation data, or other information related to the network stack. In these cases, you must use the “#” operator to retrieve this information. The table that follows provides examples of filters that access message properties and annotations.

Table 20. Accessing Message Properties and Annotations in Filter Expressions

EXPRESSION	MEANING
`HTTP.Host == "www.bing.com"`	This filter returns all HTTP messages that have a **Host** property defined with a value equal to "www.bing.com".
`Etw::ProcessId==4`	This filter, based on the **ProcessId Global Property** (see the **Field Chooser Tool Window**), returns each message in the **Analysis Grid** viewer that has a **ProcessId** that is equal to 4.
`#MessageNumber == 5`	This filter, based on the **MessageNumber Global Annotation** (see the **Field Chooser Tool Window**), returns the message in the **Analysis Grid** viewer that has a **MessageNumber** equal to 5. Note that all messages have a **MessageNumber** annotation that is defined in OPN.

Referencing Enumerations

You can refer to enumeration values by using the friendly enumeration name or by using actual field values. The table that follows depicts two different ways to refer to an enumeration.

Table 21. Using Enumerations in Filter Expressions

EXPRESSION	MEANING
`Ipv4.Protocol == IANA.ProtocolType.UDP`	This filter returns all messages with a Protocol field that is set to UDP and uses the friendly enumeration name to retrieve that value.
`Ipv4.Protocol == 17`	This filter is a variant of the previous one, in that "17" is the enumeration integer value for UDP. You should therefore expect this filter expression to return all messages that have a Protocol field that is set to UDP.

More Information

To learn more about Filter Expressions that use enumerations, including additional examples, see the Diagnosis Category topic in Filtering Live Trace Session Results.

Using Special Filtering Functions

The Filtering Language also provides a set of special functions that enable you to search for a string or hexadecimal value within a message, without having to specify an associated field name to retrieve those values. These functions are described in the following table.

Table 22. Searching for Raw Data

EXPRESSION	MEANING
`contains "some string"`	This filter returns all messages that contain "some string". The evaluation is case-insensitive and compares all supported string encodings; however, the default encoding is ASCII.
`contains "some string" caseSensitive`	This filter is similar to the first filter in this table, but here case sensitivity is applied to the string evaluation.

EXPRESSION	MEANING
`contains “some string” encoding ASCII`	This filter is similar to the first filter only the string evaluation is restricted to a specified encoding, which includes any one of the following: - ASCII - UTF7 - UTF8 - Unicode - UTF32 - BigEndianUnicode - Base64 **Note:** If you do not specify an encoding, all encodings are included in filter scope.
`contains $[AA34]`	This filter returns all messages that contain the specified binary value in hexadecimal format.

Browsing Message Origins

You can use symbols that resemble XPath notation to traverse the protocol stack of any protocol message origin tree in the following ways:

- **Directory tree traversing** — you can use the backslash symbol ("\") to traverse the protocol stack similar to the way you navigate a directory tree.
- **Origins tree traversing** — you can use the double backslash symbol ("\\") to look down the stack one or more levels.

The table that follows provides some examples of using these symbols in filter expressions to traverse the protocol stack.

Table 23. Traversing the Protocol Stack

EXPRESSION	MEANING
`\TCP\IPv4`	This filter returns top-level messages that are TCP where an IPv4 message is one level below TCP in the origins tree. Top-level is indicated by the leading “\” character.
`\HTTP\\IPv4`	This filter returns top-level messages that are HTTP where an IPv4 message exists one or more levels below HTTP in the origins tree.
`\HTTP\TCP\\Ethernet`	This filter returns top-level HTTP messages where a TCP message is directly below HTTP and somewhere below that is an Ethernet message in the origins tree.
`TCP\IPv4`	This filter returns all TCP messages from any level in the origins tree where an IPv4 message is one level below TCP. **Note:** In this filter expression, note that the leading “\” character is missing. This means that an explicit level in the origins tree is not specified. As a result, a double backslash ("\\") is implied when the leading slash (“\”) is not written in the expression, which therefore makes "TCP" semantically equivalent to "\\TCP".

You can also combine traversing notation with various field values and operators, as indicated in the table that follows.

Table 24. Combining Traversing with Field Values and Operators

EXPRESSION	MEANING
`\HTTP\TCP.Port == 80\IPv6 || UDP`	This filter returns all top-level HTTP messages with a TCP transport where the Port (SourcePort or DestinationPort) is equal to 80, and with either an (IPv6 message directly below) OR (UDP message at any level below) in the origins tree.
`\TCP.Port == 80\IPv4.Address == $[C0A80101]`	This filter returns all top-level TCP messages that have a Port value equal to 80, with IPv4 messages below that have an Address value equal to 192.168.1.1.

Using Aliases

In OPN, protocol fields can be grouped under a common name or alias by using an OPN aspect. The protocols that are included with Message Analyzer use this aspect to declare groups or "aliases" for commonly used fields. For example, the "Address" alias for IPv4 messages is defined to include SourceAddress and DestinationAddress fields, which makes it possible to apply a filter such as `*Address == 10.0.1.13`. This particular filter expression returns all messages that have either SourceAddress or DestinationAddress set to this IP address value, as it is semantically equivalent to the expression `*SourceAddress == 10.0.1.13 || *DestinationAddress == 10.0.1.13`. When fields are grouped under a common name such as "Address", it implies the disjunction of each of the components in the group. Another example of a common alias is "Port", which represents the SourcePort and DestinationPort fields. This is why applying a filter expression such as `*Port==443` will return messages from any protocol that defines a SourcePort or DestinationPort field that has a value of 443.

> **IMPORTANT**
> As related to OPN aspects, aliases are not actual message fields and therefore you cannot display an *alias* column in the **Analysis Grid** viewer column layout.

However, this should not be confused with the Message Analyzer **Aliases** feature that enables you to create and substitute a friendly name for a field value with a cryptic name that is difficult to keep track of. For more information about this feature, see Using and Managing Message Analyzer Aliases.

> **NOTE**
> OPN aspects perform a similar function as attributes do in the C# programming language. For example, attributes consists of metadata that can extend the language or declarative information that a program can use at runtime.

A compilation of all grouping aliases for protocols that are provided with Message Analyzer will be available in the near future.

More Information
To learn more about OPN aspects, see the OPN Programming Guide document.

Procedures: Using the Data Filtering Features

61 minutes to read

The procedures in this section encapsulate many of the filtering functions described in the Filtering Message Data section. For each procedure, background scenario information is provided to set the context for the problem that each procedure will help to resolve. The background information describes at a high-level some common issues that you might encounter, while the procedures demonstrate how you can use Message Analyzer to isolate relevant data with the use of filters and simplify your data assessment and problem resolving process.

> **IMPORTANT**
>
> Although these procedures demonstrate the use of Message Analyzer capabilities with respect to issues specific to the HTTP, TCP, and SMB protocols, these are only a sampling of what you can accomplish with Message Analyzer, given that you can also apply the same methodologies described here to troubleshoot a wide range of other types of network traffic and events.

> **NOTE**
>
> As the person running the procedures in the examples of this section, it is assumed that you are a Network Administrator, or that you are at least familiar with networking concepts.

Procedure Overviews

A brief description of each procedure is included here for review, as follows.

Filtering a Data Retrieval Session — provides examples of the following:

- How to apply a **Session Filter** to data that is loaded into Message Analyzer through a Data Retrieval Session that targets a saved trace file as input, so you can isolate and analyze HTTP request messages sent from a specified client IP address to a poorly performing web server along with the response messages that were issued back to the client. Some HTTP analysis techniques are also provided in the procedure to help you discover possible reasons for the poor performance of a hypothetical web server.
- How to apply a **Time Filter** to data that is loaded into Message Analyzer through a Data Retrieval Session that targets input files that consolidate several message sources, so that you can view data in specific windows of time in which problems are suspected to have occurred.

More Information

To learn more about the types of input files to which you can apply a **Time Filter**, see Applying an Input Time Filter to a Data Retrieval Session.

Filtering Live Trace Session Data — provides examples of the following:

- How to collect data from a specific network interface in a Live Trace Session that uses the **Local Network Interfaces Trace Scenario**, by selecting an **Adapter** from which to capture data. This can be useful to isolate data on a specific adapter when network adapters are on different networks or load balanced on the same subnet where traffic spans multiple adapters.
- How to specify a **WFP Layer Set** filter in a Live Trace Session that uses the **Network Tunnel Traffic and Unencrypted IPSEC Trace Scenario** to capture inbound TCP traffic only, in order to limit the number of captured messages while focusing on TCP ACK traffic for troubleshooting purposes.

- How to specify a **Session Filter** in a Live Trace Session that uses the **Loopback and Unencrypted IPSEC Trace Scenario** to isolate TCP messages, in an attempt to correlate slow file transfers or sluggish web page loading performance with TCP retransmit issues.
- How to specify a **Fast Filter** that filters messages at the provider level based on a port or client IP address. The filtering is applied to a busy server where the responses are slow, in a Live Trace Session that uses the **Loopback and Unencrypted IPSEC Trace Scenario**, in order to reduce captured message volume and lower the impact on server performance while still capturing sufficient data for analysis. Demonstrates the efficiency and performance advantages of **Fast Filters** that enable you to capture the least amount of data possible to resolve a problem.
- How to specify a **Hostname** and/or **Port** filter in a Live Trace Session that uses the **Pre-Encryption for HTTPS Trace Scenario** to reduce captured traffic volume and isolate client messages sent to a specified web server that is very busy and is causing users to experience connection issues. Applies the specified filters to gather sufficient data for the assessment of possible connection problems, without imposing a high-volume data capture on a client computer that is overwhelmed with traffic. Assesses possible causes by examining HTTP **StatusCode** indicators for connection issues and errors.

Filtering Live Trace Session Results — provides examples of the following:

- How to apply an Address and Port view **Filter** to the results of a Live Trace Session that used the **Local Network Interfaces Trace Scenario**, and to thereby remove unwanted lower-layer traffic and streamline the detection of virus signatures.
- How to apply an HTTP view **Filter** to the results of a Live Trace Session that used the **Loopback and Unencrypted IPSEC Trace Scenario**, to isolate all HTTP traffic from a particular client computer, including all the origins messages in the underlying stack where a problem might occur in supporting HTTP operations. Alternatively uses the **HTTP Viewpoint** to present data from the perspective of the HTTP protocol. Demonstrates the superior capabilities of Message Analyzer compared to its predecessor Network Monitor, which returns only HTTP headers in a similar scenario.
- How to apply TCP view **Filters** to the results of a Live Trace Session that used the **Loopback and Unencrypted IPSEC Trace Scenario** to help expose the causes of TCP connection and data transmission issues such as lost TCP segments, slow data transmission rates, or broken TCP three-way handshakes.
- How to apply an HTTP filter to the results of a Live Trace Session that used the **Pre-Encryption for HTTPS Trace Scenario**, so you can isolate specific HTTP request/response messages and analyze HTTP **StatusCodes** and **Reason** phrases for evidence of a poorly performing web server.

Filtering with Color Rules and Exposing Diagnostics for TCP and SMB — how to apply **Color Rules** to the results of a Live Trace Session that used the **Network Tunnel Traffic and Unencrypted IPSEC Trace Scenario** with a **Session Filter** to expose TCP and SMB diagnostic messages, in order to prompt further analysis of faulty or erratic file share access.

> **NOTE**
>
> In the procedures of this section, placeholders within angle brackets (<>) refer to values that you enter that are specific to your system/s. However, do not include the angle brackets in your entries when testing these procedures.

IMPORTANT

If you have not logged off Windows after the first installation of Message Analyzer, please log off and then log back on before performing these procedures. This action ensures that in all subsequent logons following installation, your security token will be updated with the required security credentials from the Message Capture Users Group (MCUG). Otherwise, you will be unable to capture network traffic in **Trace Scenarios** that use the **Microsoft-PEF-NDIS-PacketCapture** provider, **Microsoft-Windows-NDIS-PacketCapture** provider, or the **Microsoft-PEF-WFP-MessageProvider**, unless you start Message Analyzer with the right-click **Run as administrator** option.

Even if you log off your system, log back on, and receive the required security credentials from the MCUG, you will still need to use the **Run as administrator** option if you want to capture message traffic in Message Analyzer **Trace Scenarios** that use the **Microsoft-Windows-NDIS-PacketCapture** provider or the **Microsoft-PEF-WFP-MessageProvider**, which both have remote capabilities. Because of the inherent remote capabilities of these message providers, additional security restrictions must be applied.

Filtering a Data Retrieval Session

This section provides examples of possible ways to filter data that is loaded into Message Analyzer from saved trace or log files. The examples include the use of an HTTP **Session Filter** that is applied to a saved trace file and an adjustable **Time Filter** that is applied to one or more log files.

Applying a Session Filter to Saved Trace Data

The hypothetical high-level issue in this example is that a Network Administrator has client browsers that send requests to a web server that is being constrained, possibly by high message volumes, connection problems, or other issues. The administrator loads data from one or more saved trace files containing messages that were originally captured on one or more client computers by using the **Loopback and Unencrypted IPSEC Trace Scenario**, which reduces most network noise at the Data Link Layer. The example assumes that the trace ran over an adequate time period to reproduce the connection issues that clients are experiencing. Messages are loaded into Message Analyzer through a Data Retrieval Session with a **Session Filter** applied that isolates HTTP "GET" requests (displayed under HTTP operation nodes in the **Analysis Grid** viewer) that were sent to a specified web server from a particular client IP address.

The purpose of the **Session Filter** in this case is to enable the administrator to isolate HTTP operations and facilitate analysis of the HTTP request and response messages that were interchanged between a specified web server and a specified client IP address, which resulted in possible transport errors, request timeouts, or other faulty status indications. In this example, the administrator also discovers possible reasons why the web server performed poorly by examining the HTTP **StatusCode** indicators and **ReasonPhrases** that were sent back to the client computer in HTTP response messages.

NOTE

A **Session Filter** enables you to isolate specific messages and limit the amount of data stored in memory, which provides a performance advantage, but can be more expensive in terms of processing time. A view **Filter** is somewhat more efficient and makes it easier to work with your data after it displays in Message Analyzer.

For example, a view **Filter** can be applied and then removed because all of your original data is preserved in the Message Store, whereas with a **Session Filter**, all messages that are not targeted to pass the filter are removed before any data is displayed and can only be retrieved by re-running the Data Retrieval Session with the **Session Filter** removed. If you want to re-run a session, click the **Edit Session** button on the global Message Analyzer toolbar to display the **Edit Session** dialog, from where you can reconfigure the session and run it again, as described in Editing Existing Sessions.

To apply a Session Filter to loaded data for HTTP message analysis

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer**

icon to launch Message Analyzer. If you have not logged off and back on after first installing Message Analyzer, then start Message Analyzer with the right-click **Run as Administrator** option.

2. On the global Message Analyzer toolbar, click the **New Session** button to display the **New Session** dialog.
3. Under **Add Data Source** in the **New Session** dialog, click the **Files** button to display the **Files** tab along with the associated session configuration features that it contains in the **New Session** dialog.
4. On the toolbar of the **Files** tab in the **New Session** dialog, click **Add Files** to display the **Open** dialog and then navigate to the trace file/s containing the saved trace data you want to work with.
5. In the **Open** dialog, select the file/s containing the data you want to load into Message Analyzer, and then click **Open**.

   The files list is populated with the file/s you selected.
6. In the files list, ensure that there is a check mark in the check box next to the file/s containing the data you want to load into Message Analyzer.
7. In the **Session Filter** text box of the **New Session** dialog, enter the following Filter Expression and substitute appropriate values for the placeholder italic values. In this expression, note that the value for the "Source" phrase is the client IP address:

   ```
   HTTP.Request.Method=="GET" && *HTTPHost == "www.hostname.com" && *Source == <192.168.1.1>
   ```
8. Verify that the **Analysis Grid** viewer is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.
9. Click the **Start** button in the **New Session** dialog to begin loading data into Message Analyzer from the specified file/s.

   The loaded data displays in the **Analysis Grid** viewer.
10. On the **Analysis Grid** viewer toolbar, click the **Add Columns** button to display the **Field Chooser Tool Window**.
11. In **Field Chooser**, double-click the **StatusCode** and **ReasonPhrase** fields for **Response** messages in the **HTTP** message hierarchy to add **StatusCode** and **ReasonPhrase** columns to the **Analysis Grid** viewer column layout.
12. Right-click the **StatusCode** column and select the **Group** command to isolate the status data into separate groups with identical status codes for ease of analysis.

> **TIP**
> You can right-click the **StatusCode** label and select the **Expand All Groups** menu item to expand all the nodes of grouped data.

13. In the **Analysis Grid** viewer, examine the **StatusCode** and **Reason Phrase** values to determine any problem areas that might indicate a poorly performing web server, as described in the status code table in the Addendum 2: HTTP Status Codes section of this documentation.

NOTE

You can also review the default **TimeElapsed** column in the **Analysis Grid** viewer to verify the elapsed time for HTTP "GET" methods, or for entire HTTP operations to complete. High values for elapsed time in the latter "operations" case may be an indication of network latency or TCP retransmit issues due to dropped packets.

You might also add a **ResponseTime** column (from the **Global Annotation** category in **Field Chooser**) to examine the server response times to request messages. This can provide an indication of whether or not a server is performing slowly, as it measures the time difference between the last time-stamped message in the request stack and the **Timestamp** of the first response message that is sent back to the requesting node. Note that you can create a similar configuration for any protocol that makes use of request/response pairs.

14. Expand the top-level HTTP operations nodes and child nodes in the **Analysis Grid** viewer to expose the origins tree for HTTP messages to determine whether any TCP diagnosis errors are present. You can do this by clicking the **Diagnostics** icon in the **DiagnosisTypes** column in the **Analysis Grid** viewer for any message to display diagnosis error details inline. Alternatively, you can view diagnostic information more effectively by doing any of the following:

    - Open the **Diagnostics Tool Window** by selecting it in the **Windows** submenu of the global Message Analyzer **Tools** menu, to display a summary of diagnosis messages for your data set. You can then select diagnosis messages in the **Diagnostics** window to drive the interactive highlighting of corresponding messages in the **Analysis Grid** viewer for further examination of diagnosis errors in those messages.
    - Perform the **Group** command on the **Diagnosis** column in the **Analysis Grid** viewer to isolate errors by type, by right-clicking the **Diagnosis** column and selecting the **Group** command.
    - Click the **Diagnosis** column in the **Analysis Grid** viewer until you sort and bubble up any diagnosis errors that might have occurred.

TIP

For ease of analysis, you might also consider displaying only HTTP messages encapsulated in top-level operations with no layers above, by applying an **HTTP Viewpoint** from the **Viewpoints** drop-down list on the Filtering toolbar that appears above the **Analysis Grid** viewer.

NOTE

Diagnosis messages that specify lost TCP segments and retransmits might be an indication of packets being dropped by the network, or could indicate TCP performance issues related to TCP **Window** size and/or **WindowsScaleFactor** settings.

Applying a Time Filter to Loaded Log Data

The hypothetical high-level issue in this example is that a Network Administrator has a large volume of data that was collected in one or more log files and he or she wants to view the data in a specific window of time where failures are suspected to have occurred. In this scenario, the administrator will consolidate one or more related log files as input to a Data Retrieval Session and will apply a **Time Filter** that defines a time window in which to view messages. By using the Message Analyzer **Time Filter**, the administrator will limit the amount of data being loaded, reduce the loading time, and as a result, realize better performance. Optionally, the administrator can apply a **Session Filter** to the input data to isolate messages to a specific client IP address, if the log file format supports IP addresses, thereby reducing message count and helping to streamline the message analysis process.

NOTE

Some log files, such as text-based log files in proprietary format, may require you to create an OPN configuration file so that the log messages can be fully parsed by Message Analyzer, as described in Working With Special Input Requirements. However, note that you can select from a number of built-in **Text Log Configuration** files on the toolbar of the **Files** tab in the **New Session** dialog, during Data Retrieval Session configuration. If your log file cannot be parsed by one of the built-in configuration files, you will need to create one, as described in Opening Text Log Files.

To apply a Time Filter to input file data and view results in a selected time window

1. From the **Start** menu, **Start** page, or task bar of your computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer. If you have not logged off and back on after first installing Message Analyzer, then start Message Analyzer with the right-click **Run as Administrator** option.
2. Click the **New Session** button on the Message Analyzer **Start Page** to display the **New Session** dialog.
3. Under **Add Data Source** in the **New Session** dialog, click the **Files** button to display the **Files** tab along with the associated session configuration features that it contains in the **New Session** dialog.
4. On the toolbar of the **Files** tab in the **New Session** dialog, click **Add Files** to display the **Open** dialog and then navigate to the trace file/s containing the saved trace data you want to work with.
5. In the **Open** dialog, select the file/s containing the data you want to load into Message Analyzer, and then click **Open**.

   The files list is populated with the file/s you selected.
6. In the files list, ensure that there is a check mark in the check box next to the file/s containing the data you want to load into Message Analyzer.

   After you select the files to include in the Data Retrieval Session, the **Time Filter** pane on the **Files** tab will be populated with **Start Time** and **End Time** values that are derived from the input file/s, in addition to the **Total Messages** count. Note that the time values that initially display create a window that is inclusive of the earliest and latest time values that Message Analyzer detects from the currently selected files in the input files list.
7. Select the **Use Start Filter** and **Use End Filter** check boxes in the **Time Filter** pane and then adjust the **Time Filter** slider controls to set the time window that contains the data you want to examine.

   The **Start Time** and **End Time** values track the position of the slider controls and the **Filtered Messages** count changes to indicate the number of messages that the Data Retrieval Session will load from the selected log files, for example, one or more *.etl files, based on the current **Time Filter** settings.

NOTE

If the **Start Time** and **End Time** values do not display in the **Time Filter** pane for your log, you can manually specify starting and ending date-times in a format that is appropriate for your log file/s. Thereafter, time window adjustments should track in accordance with the format that you specified. Also be aware that **Start Time** and **End Time** values will appear for *.log files only after they are parsed, meaning that you can only apply a **Time Filter** to messages in these file types through the **Edit Session** dialog, unless the configuration file specifies the time stamp format.

8. Optionally, enter the following Filter Expression in the **Session Filter** text box of the Data Retrieval Session configuration, while substituting appropriately for the value in italics, so that you can isolate messages from a particular client IP address — providing that this action is appropriate for the log files you are working with as indicated earlier:

```
*Source==<192.168.1.1>
```

9. Verify that the **Analysis Grid** viewer is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.

10. Click the **Start** button to begin loading data from the specified input log file/s.

    The loaded and filtered data displays in the **Analysis Grid** viewer to enable you to examine the messages in the time window that you specified.

11. To optionally load log data in different windows of time, click the **Edit Session** button on the global Message Analyzer toolbar to open the **Edit Session** dialog, from where you can readjust the slider controls to configure a different **Time Filter** window. Note that you will need to enter the **Full Edit** mode in the **Edit Session** dialog to enable the controls.

> **NOTE**
> If you alter the original Data Retrieval Session configuration after entering the **Full Edit** mode, by making changes other than adding or removing files, Message Analyzer will perform a required reload of all data.

12. When reconfiguration of the time window in your Data Retrieval Session is complete, click the **Apply** button to start loading data based on the new **Time Filter** or other session reconfiguration.

> **NOTE**
> After you apply a **Time Filter** to an input data files configuration, all messages that are outside the time window you specified are removed from displayed results. If you want to restore all messages from your log files, you will need to reload all data with the **Time Filter** removed.

However, as an alternative to applying an input **Time Filter** prior to loading data from saved files, you can simply load all input log data as is and then apply a **Time Filter** to session results from the Filtering toolbar to isolate data to a specified time window. The action of the latter **Time Filter** can be quickly undone by simply clicking the **Remove** button on the Time Filter panel of the Filtering toolbar, so you can conveniently return to the original data set to review all messages or apply different time window filtering as required, without having to reload any data. However, there are some factors to consider when choosing to use an input **Time Filter** versus a session results **Time Filter**, as described in Applying a Time Filter to Session Results. For example, when loading data from large log files without an input **Time Filter**, Message Analyzer performance may diminish.

Filtering Live Trace Session Data

This section provides examples of possible ways to filter data as it is being collected from several Live Trace Sessions that separately utilize one of the following default **Trace Scenarios**:

- **Local Network Interfaces**
- **Network Tunnel Traffic and Unencrypted IPSEC**
- **Loopback and Unencrypted IPSEC**
- **Pre-Encryption for HTTPS**

The types of filters that are applied in these four scenarios, in order, consist of an **Adapter** filter, **WFP Layer Set** filter, **Session Filter**, **Fast Filter**, and a **HostnameFilter** and **PortFilter**, as described in the sections that follow.

Selecting a Network Adapter to Filter a Local Network Interfaces Trace
The hypothetical high-level issue in this example is that a Network Administrator has a computer that is acting as a firewall with two adapters on different networks, and he or she wants to look at traffic on one adapter only,

possibly for evidence of packets dropped at the NDIS layers. Message Analyzer enables the administrator to specify a particular adapter on which to capture data in a **Local Network Interfaces Trace Scenario**, by selecting the adapter in the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog that is accessible from the **New Session** dialog during session configuration, while deselecting all other adapters. A trace at the Data Link Layer ensures that the administrator can capture relevant data for the selected adapter.

> **IMPORTANT**
>
> If you are running the **Local Network Interfaces** scenario on a computer with the Windows 7, Windows 8, or Windows Server 2012 operating system, you will be working with the **Microsoft-PEF-NDIS-PacketCapture** provider. However, if you are running this scenario on a computer with the Windows 8.1, Windows Server 2012 R2, or Windows 10 operating system, you will be working with the **Microsoft-Windows-NDIS-PacketCapture** provider.
>
> Please note that the **Advanced Settings** dialogs for these two providers are very different and contain unique filtering configuration capabilities. The procedure that follows focuses on the **Microsoft-PEF-NDIS-PacketCapture** provider and its configuration capabilities. However, to configure adapter selection in any **Trace Scenario** that uses the **Microsoft-Windows-NDIS-PacketCapture** provider, refer to the Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog topic for further details.

To filter data on a specific network interface in a Local Network Interfaces trace

1. From the **Start** menu, **Start** page, or task bar of the target computer, click the **Microsoft Message Analyzer** icon to launch Message Analyzer. If you have not logged off and back on after first installing Message Analyzer, then start Message Analyzer with the right-click **Run as Administrator** option.
2. Click the global Message Analyzer **File** menu, click **New Session**, and then select **Blank Session** in the **New Session** submenu to display the **New Session** dialog.
3. Under **Add Data Source** in the **New Session** dialog, click the **Live Trace** button to display the **Live Trace** tab along with the associated session configuration features that it contains in the **New Session** dialog.
4. From the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, select the **Local Network Interfaces Trace Scenario**.

   The **ETW Providers** list is populated with the **Name** and **Id** (GUID) of the **Microsoft-PEF-NDIS-PacketCapture** (or **Microsoft-Windows-NDIS-PacketCapture**) provider, along with the **Configure** link, which provides access to the provider **Advanced Settings** dialog.
5. In the **ETW Providers** list on the **Live Trace** tab, click the **Configure** link next to the **Microsoft-PEF-NDIS-PacketCapture** provider **Id** to display the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, as shown in Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.
6. In the **Advanced Settings** dialog, click the **Provider** tab and then deselect the **In** and **Out** check boxes for the **Machine** node in the **System Network** tree grid to deselect the **In** and **Out** check boxes for all adapters.
7. In the **System Network** tree grid, select the adapter on which to capture data by selecting the **In** and **Out** check boxes for that particular adapter. Ensure that all other adapters are unselected and then click **OK** to exit the dialog.

   The **Microsoft-PEF-NDIS-PacketCapture** provider is now set to capture traffic in both directions on the network interface that you specified.

> **NOTE**
>
> You have the option to select only the **In** or **Out** check boxes for any adapter in the **System Network** tree grid, so that you can monitor traffic in a specified direction only. However, in most cases, it is prudent to monitor traffic in both directions. Note that you also have the option to specify a **Fast Filter** for this Live Trace Session in the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, to focus on specific results and improve performance.

8. Verify that the **Analysis Grid** is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.
9. Click the **Start** button in the **New Session** dialog to begin capturing data.

   The captured messages start to accumulate in the Message Analyzer **Analysis Grid** viewer.
10. While Message Analyzer is capturing data, attempt to reproduce any issue that is related to the reason for capturing data on a particular adapter, although this might not be necessary if the issue is an NDIS layer that is simply dropping packets.
11. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.
12. Examine the data you captured, as appropriate.

> **TIP**
>
> When using the **Microsoft-Windows-NDIS-PacketCapture** provider in the **Local Network Interfaces** scenario, you can configure filtering that can help determine whether an NDIS layer is dropping packets.
>
> For example, by selecting the **All Layers** check box in the **Advanced Settings – Microsoft-Windows-NDIS-PacketCapture** dialog, you can specify that packets are intercepted on all layers of the NDIS stack. In addition, by selecting both the **Ingress** and **Egress** check boxes in the dialog, you can specify the traversal path (direction up and down the stack) in which packets are intercepted. By specifying all layers and both traversal paths for an adapter that is dropping packets, you make certain that dropped packet events will be generated for any layer. If packets are being dropped, you should be able to expose them in Message Analyzer as ETW events that have certain characteristics, as described in step 24 of the procedure Capture Traffic on a Remote Host.

More Information

To learn more about configuring the **Microsoft-PEF-NDIS-PacketCapture** provider from the **Advanced Settings - Microsoft-PEF-NDIS-PacketCapture** dialog, including how to assign **Fast Filter Groups** to selected adapters, see Using the Advanced Settings - Microsoft-PEF-NDIS-PacketCapture Dialog.

To learn more about the unique filtering configurations and other advanced capabilities that are available for the **Microsoft-Windows-NDIS-PacketCapture** provider, including remote tracing, see Using the Advanced Settings - Microsoft-Windows-NDIS-PacketCapture Dialog.

Applying a WFP Layer Set Filter to a Network Tunnel Traffic and Unencrypted IPSEC Trace

The hypothetical high-level issue in this example is that a Network Administrator wants to run a Live Trace Session that directionally isolates TCP traffic on a client computer to detect connectivity issues while at the same time reduce the amount of TCP traffic that is captured to streamline performance. Message Analyzer enables the administrator to use the **WFP Layer Set** filter configuration of the **Microsoft-PEF-WFP-MessageProvider** in the **Network Tunnel Traffic and Unencrypted IPSEC Trace Scenario** to isolate inbound V4 traffic at the Transport Layer as a possible starting point for detecting network connectivity issues. Note that the **WFP Layer Set** filter configuration also enables capture of outbound V4 traffic at the Transport Layer in addition to inbound and outbound V6 traffic, in any combination.

To apply WFP Layer Set filtering to a Network Tunnel Traffic and Unencrypted IPSEC trace and isolate TCP diagnosis traffic

1. On the client computer, perform steps 1 through 3 of the previous procedure: To filter data on a specific

network interface in a Local Network Interfaces trace.

2. From the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, select the **Network Tunnel Traffic and Unencrypted IPSEC Trace Scenario**.

   The **ETW Providers** list is populated with the **Name** and **Id** (GUID) of the **Microsoft-PEF-WFP-MessageProvider**, along with the **Configure** link, which provides access to the **Advanced Settings** dialog for this provider.

3. In the **ETW Providers** list on the **Live Trace** tab, click the **Configure** link next to the **Id** of the **Microsoft-PEF-WFP-MessageProvider** to display the **Advanced Settings - Microsoft-PEF-WFP-MessageProvider** dialog.

4. In the **Advanced Settings** dialog, click the **Provider** tab to display the **WFP Layer Set** inbound and outbound Transport Layer filters.

   By default, the **WFP Layer Set** filter configuration is set to capture messages for all inbound and outbound transports.

5. In the **WFP Layer Set** pane on the **Provider** tab, deselect the **Outbound Transport V4**, **Inbound Transport V6**, and **Outbound Transport V6** check boxes; then ensure that only the **Inbound Transport V4** check box is selected.

   With this **WFP Layer Set** configuration, the TCP messages that this Live Trace Session captures will be inbound TCP packets only.

6. Verify that the **Analysis Grid** is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.

7. Click the **Start** button to begin capturing data.

   The captured messages start to accumulate in the Message Analyzer **Analysis Grid** viewer.

8. While Message Analyzer is capturing data, attempt to reproduce the conditions that result in network connectivity issues being experienced by the client.

9. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.

10. From the **Viewpoints** drop-down list on the Filtering toolbar above the **Analysis Grid** viewer, select the **TCP** item to filter all TCP traffic to top-level so you can more easily examine the TCP inbound traffic that you captured for signs of connection or transmission issues.

    For example, with the **Inbound Transport V4** filter, you could focus on TCP acknowledgement (ACK) traffic only to see whether you have a large number of duplicate ACK messages. You could also add the **TCPDiagnosis** column (see the TCP message hierarchy in **Field Chooser**) to the **Analysis Grid** viewer as a new data column to see diagnosis message summaries that provide useful diagnostic information. You might also execute the right-click **Group** command on the new column to better organize the diagnosis data.

    Other measures you can take to expose connectivity or transmission issues include applying either of the following view **Filters** in the text box of the Filtering toolbar above the **Analysis Grid**: `*TCPDiagnosis contains "Dup-Ack"` or `*TCPDiagnosis contains "retransmitted"`.

> **TIP**
> You can also specify these filters in separate Filter panels. The first Filter panel displays by default, where you can enter the first filter in the Filter Expression text box of that panel. To add the second filter to the text box of another Filter panel, select **Add Filter** from the **Add Filter** drop-down list on the Filtering toolbar to display the second Filter panel with its Filter Expression text box and enter the text of the second filter into it. You can then apply and remove these filters independently by clicking the respective **Apply** and **Remove** buttons on the Filter panels.

You might also reconfigure the **WFP Layer Set** filtering by clicking **Edit Session** on the global Message Analyzer toolbar, clicking the **Configure** link for the **Microsoft-PEF-WFP-MessageProvider** in the **ETW Providers** list to open the **Advanced Settings** dialog, and then selecting the **Outbound Transport V4** filter only on the **Provider** tab of the dialog. Thereafter in the trace results, apply the following view **Filter** so you can focus on outbound TCP synchronization traffic from the client: `tcp.syn==true`.

In any case, you can also look at diagnosis error messages in the default **DiagnosisTypes** column of the **Analysis Grid** viewer to see what TCP issues you might have. The goal in creating these **WFP Layer Set** filtering configurations is to capture the least amount of traffic to resolve TCP connectivity and/or transmission problems.

More Information
To learn more about Advanced Settings for the **Microsoft-PEF-WFP-MessageProvider**, see Using the Advanced Settings- Microsoft-PEF-WFP-MessageProvider Dialog.

Applying a Session Filter to a Loopback and Unencrypted IPSEC Trace
The hypothetical high-level issue in this example is that a Network Administrator has a client computer that is experiencing slow file transfer activity or slowly loading web pages. A high number of TCP retransmits could be responsible for the delays. This could be symptomatic of inappropriate TCP **Window** size and **WindowsScaleFactor** settings, or possibly packets are being dropped by the network or client firewall.

> **NOTE**
> Retransmits can also be the result of the memory capacity of a router that cannot keep up with the traffic.

In this example, the administrator runs a trace on the client computer for a significant time period to gather data, while using the **Microsoft-PEF-WFP-MessageProvider** in a **Loopback and Unencrypted IPSEC Trace Scenario** to minimize low-level traffic from the Data Link Layer and a **Session Filter** to capture only top-level TCP messages or messages containing TCP in the stack. When the trace is complete, the administrator then adds **Window**, **Options**, and **TCPDiagnosis** columns to the **Analysis Grid** viewer column **Layout** and also applies a diagnostic view **Filter** to isolate and review any **Warning** level diagnosis messages that indicate retransmits occurred. The administrator also uses the **Group** command in the **Analysis Grid** viewer to gather the data into a convenient format that quickly consolidates and exposes related data of interest in groups.

If there are a significant number of TCP retransmits occurring, the administrator can do the following:

- Ensure that loopback traffic is being filtered out by applying **IPv4** and **IPv6 Fast Filters** such as `!=127.0.0.1` and `!=::1`, respectively, or filter out the IP address that a local application is using, if applicable. Otherwise, there can be duplication of TCP retransmit messages. You can configure these **Fast Filters** in the **Advanced Settings - Microsoft-PEF-WFP-MessageProvider** dialog. Note that the **Local Loopback Network Trace Scenario** sets these filters by default to remove loopback traffic.
- Review the TCP **Window** sizes of TCP segments in the **Windows** column of the **Analysis Grid** viewer and observe whether the values vary significantly. To examine messages for the client IP address only, the administrator can apply the following filter, while substituting appropriately for the IP address in italics:

```
*SourceAddress==<192.168.1.1>
```

- Review TCP **Options** to ensure that the **WindowsScaleFactor** and **Window** size are set to reasonable values, for example, WindowsScaleFactor=2 and Window size=64k.
- Verify whether the network is dropping packets, for example, a router with excessive memory utilization might be unable to keep pace with packet traffic.
- Following the initial trace results, verify whether client computer Firewall rules are causing TCP packets to be dropped, or whether the **PEF-WFP-MessageProvider** is causing it; see *Checking for Dropped Packets* in the next procedure.

To apply a Session Filter to a Loopback and Unencrypted IPSEC trace and evaluate TCP diagnostics

1. On the client computer, perform steps 1 through 3 of the previous procedure: To filter data on a specific network interface in a Local Network Interfaces trace.
2. From the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, select the **Loopback and Unencrypted IPSEC Trace Scenario**.

   The **ETW Providers** list is populated with the **Name** and **Id** (GUID) of the **Microsoft-PEF-WFP-MessageProvider**, along with the **Configure** link, which provides access to the provider **Advanced Settings** dialog.
3. In the **Session Filter** text box of the **New Session** dialog, enter the following **Session Filter** to capture top-level TCP messages only, where TCP messages are not in the message (origins) stack:

   ```
   \TCP
   ```

   > **NOTE**
   > By using the **Loopback and Unencrypted IPSEC Trace Scenario** with the **Microsoft-PEF-WFP-MessageProvider** and applying the specified **Session Filter**, the collected message count on the client computer will be reduced.
4. Verify that the **Analysis Grid** viewer is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.
5. Click the **Start** button in the **New Session** dialog to begin capturing data.

   The captured messages start to accumulate in the Message Analyzer **Analysis Grid** viewer.
6. While Message Analyzer is capturing data, attempt to reproduce the conditions that cause the client to experience the indicated issues.
7. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.
8. From the **Field Chooser Tool Window**, add TCP **Window**, TCP **Options**, and **TCPDiagnosis** columns to the **Analysis Grid** viewer column layout.
9. Filter out all TCP messages that do not have a diagnosis error by entering the following Filter Expression in the text box of the default Filter panel on the Filtering toolbar and then click the **Apply** button on the toolbar:

   ```
   TCP#DiagnosisLevels
   ```
10. Optionally, add the client IP address to the diagnosis filter as follows, to isolate diagnostic messages to the client IP address only. In the Filter Expression, substitute appropriately for the italic address value:

    ```
    TCP#DiagnosisLevels && *Source == <192.168.1.1>
    ```

11. Right-click the **TCPDiagnosis** column in the **Analysis Grid** viewer and select the **Group** command to group the different diagnostic error types.

12. Scroll down through the groups until you find one or more groups labeled **Retransmitted**.

    If there are a high number of **Warning** level "Retransmitted" diagnosis messages, then notwithstanding dropped packets, you should determine whether the TCP **Window** size and the **WindowsScaleFactor** of incoming TCP messages are set to appropriate values on the client, in the steps that follow.

13. Remove the **TCP#DiagnosisLevels** filter and then apply the following view **Filter** in the text box of the default Filter panel to filter out all TCP messages that do not contain TCP **Options**:

    ```
    TCP.Options ~= nothing
    ```

> **NOTE**
>
> The above filter will return only TCP messages that contain TCP Options and all other messages will be filtered out of the display. Alternatively, if you want to view all TCP messages in relation to IP conversations, including those that have TCP **Options**, you can use the **Field Chooser** window to add a **Network** column (message hierarchy path is IPv4.Datagram.Network) and a **Transport** column (message hierarchy path is TCP.Segment.Transport) to the **Analysis Grid** viewer. You can then use the **Group** command on each of these columns to pivot the data into differently organized data displays that provide alternate contexts for viewing TCP **Options**, as further described by the procedure in To apply TCP view Filters to Loopback and Unencrypted IPSEC trace results and expose TCP diagnostics.

14. Highlight TCP messages containing TCP **Options** as necessary and review the **WindowsScaleFactor** option settings, especially for messages that have a low **Window** size value, to determine that they are set to reasonable values. The TCP messages of interest should be those that comprised the SYN request and SYN/ACK response messages of three-way handshakes, as this is where the Option settings are negotiated. You might also consider executing the **TCP Three-Way Handshake Pattern Expression** to see the context in which these messages occur. This will also enable you to obtain further information that assists in TCP troubleshooting. See the Pattern Match Viewer section for more details

Checking for Dropped Packets

If you want to perform additional checks to determine whether dropped packets are causing TCP retransmits, perform the following steps:

1. In the current Analysis Session, click the **Edit Session** button on the global Message Analyzer toolbar to open the **Edit Session** dialog.

2. In the **ETW Providers** list, click the **Configure** link for the **Microsoft-PEF-WFP-MessageProvider** to open the **Advanced Settings - Microsoft-PEF-WFP-MessageProvider** dialog.

3. On the **ETW Core** tab of the dialog, click the **Keywords(Any)** ellipsis and select the **ut:Dropped** event **Keyword** option in the **ETW Keyword Filter Property** dialog.

4. On the **Providers** tab of the dialog, select the **Select Discarded Packet Events** option.

5. Click **OK** to exit the dialog and then click the **Apply** button in the **Edit Session** dialog to apply the changes you specified.

6. After the results have stabilized, check whether the Firewall is blocking packets:

   - Click the **Show Column Filter Row** icon in the upper left corner of the **Analysis Grid** viewer to display the **Column Filter** text boxes.

   - Enter the term "Discard" in the **Column Filter** text box beneath the **Summary** column header in the **Analysis Grid** viewer to expose any messages that might indicate the Firewall was involved in blocking traffic.

7. Check whether the **Microsoft-PEF-WFP-MessageProvider** is dropping packets:

   - On the Filtering toolbar, click the **Viewpoints** drop-down list and then select the **ETW Layer Viewpoint** to display only ETW messages in the **Analysis Grid** viewer.
   - Apply the following view **Filter** to the results to verify whether the **ut:Dropped** event is present in one or more ETW messages:
     `ETW.EtwProviderMsg.EventRecord.Header.Descriptor.Keywords==0x0000010000000000`
     You can also look for the KW_DROPPED flag value in the **Details Tool Window** after selecting any ETW message.

More Information
To learn more about filtering for TCP diagnostic messages, see the procedure To apply TCP view Filters to Loopback and Unencrypted IPSEC trace results and expose TCP diagnostics.

Applying Fast Filters to a Loopback and Unencrypted IPSEC Trace
The hypothetical high-level issue in this example is that a Network Administrator has a busy server where the responses are slow, but he or she does not want to impact the server with a high-volume trace when troubleshooting issues. Message Analyzer enables the administrator to accommodate this situation and determine why the server is behaving this way, by using a **Fast Filter** based on a **Port** to narrow down the traffic captured on the server. If the problem seems to be related to a particular user, the administrator can use a **Fast Filter** based on a specified client IP address. The intent of this scenario is to demonstrate the efficiency and performance advantages that can be achieved when applying **Fast Filters**, which enable the administrator to capture the least amount of data possible to resolve a problem.

To apply Fast Filters to a Loopback and Unencrypted IPSEC trace and isolate specific messages to reduce data volume

1. On the server computer, perform steps 1 through 3 of the previous procedure: To filter data on a specific network interface in a Local Network Interfaces trace.
2. From the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, select the **Loopback and Unencrypted IPSEC Trace Scenario**.

   The **ETW Providers** list is populated with the **Name** and **Id** (GUID) of the **Microsoft-PEF-WFP-MessageProvider**, along with the **Configure** link, which provides access to the provider **Advanced Settings** dialog.
3. In the **ETW Providers** list on the **Live Trace** tab, click the **Configure** link next to the **Id** of the **Microsoft-PEF-WFP-MessageProvider** to display the **Advanced Settings - Microsoft-PEF-WFP-MessageProvider** dialog.
4. In the **Advanced Settings** dialog, click the **Provider** tab to display the **Fast Filters** configuration.

   By default, there are no **Fast Filters** set.
5. In the **Fast Filters** pane, click the **Fast Filter 1** drop-down list and select the **TCP port** item.
6. In the text box to the right of the **Fast Filter 1** drop-down list, specify a port number such as `<80>`.
7. Alternatively, to isolate the data to a specific client IP address, click the **Fast Filter 2** drop-down list and select the **IPv4** item.
8. In the text box to the right of the **Fast Filter 2** drop-down list, specify a client IPv4 address in a format similar to the following: `<192.168.1.1>`.
9. Click the **OK** button to exit the **Advanced Settings** dialog.
10. Verify that the **Analysis Grid** is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.

11. Click the **Start** button in the **New Session** dialog to begin capturing data.

    The captured messages start to accumulate in the **Analysis Grid** viewer.

12. While Message Analyzer is capturing data, attempt to reproduce the conditions that cause the server to experience the indicated issues.

13. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.

14. Examine the captured data to confirm that slow server responses is an issue.

    For example, if you are working with message data from a protocol that uses request/response pairs (viewed as Operation nodes in the **Analysis Grid** viewer) such as HTTP, SMB, or DNS; ports 80, 445, and 53, respectively; you can add a **ResponseTime** column to the **Analysis Grid** viewer and then sort that column to view the messages that have the highest server response times. Note that high server response times can typically rule out network issues as the cause of delays, providing that operation **TimeElapsed** values are reasonable.

More Information
To learn more about the server **ResponseTime** field, see Average Elapsed Time for Operations.

Applying Hostname and Port Filters to a Pre-Encryption for HTTPS Trace
The hypothetical high-level issue in this example is that a Network Administrator is dealing with a client browser that has difficulty connecting with one or more web sites. The client is already overwhelmed with network traffic, so the administrator wants to determine what might be causing the connection problems without imposing a high-volume data capture on the overwhelmed client computer.

In this example, the administrator uses the **Pre-Encryption for HTTPS Trace Scenario** with the **Microsoft-Pef-WebProxy** provider and a **Hostname Filter** to isolate messages sent to and from a particular web site, and a **Port Filter** to isolate messages on a particular port, such as 80. This scenario is useful for achieving better performance than using a comparable **Session Filter** that specifies a particular destination IP address and port number. The **Hostname Filter** and **Port Filter** act similarly to the way **Fast Filters** do in that they cause less data to be collected, allow for minimal parsing, and therefore have less impact on the computer where the trace is run. In this scenario, the administrator analyzes HTTP **StatusCode** and **ReasonPhrase** indicators that can reflect connection issues, request timeouts, elapsed time/delays, or other errors that occur when a client is attempting to connect to a specified HTTP web server that might also be very busy.

> **NOTE**
> In this scenario, you will be unable to analyze the message data for possible TCP issues because the **Microsoft-Pef-WebProxy** provider captures messages above the Transport Layer.

To apply HTTP filters to a Pre-Encryption for HTTPS trace and expose status indications

1. On the client computer, perform steps 1 through 3 of the previous procedure: To filter data on a specific network interface in a Local Network Interfaces trace.

2. From the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, select the **Pre-Encryption for HTTPS Trace Scenario**.

    The **ETW Providers** list is populated with the **Name** and **Id** (GUID) of the **Microsoft-Pef-WebProxy** provider, along with the **Configure** link, which provides access to the provider **Advanced Settings** dialog.

3. In the **ETW Providers** list on the **Live Trace** tab, click the **Configure** link next to the **Id** of the **Microsoft-Pef-WebProxy** provider to display the **Advanced Settings - Microsoft-Pef-WebProxy** provider dialog.

4. In the **Advanced Settings** dialog, click the **Provider** tab to display the filter configuration.
5. In the **Hostname Filter** text box, enter a host name value in a format similar to the following:

   ```
   www.bing.com
   ```

6. In the **Port Filter** text box, enter an HTTP port number in a format similar to the following:

   ```
   80
   ```

7. Click **OK** to exit the **Advanced Settings** dialog.
8. Verify that the **Analysis Grid** viewer is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.
9. Click the **Start** button in the **New Session** dialog to begin capturing data.
10. While Message Analyzer is running, attempt to reproduce the conditions that cause the connection issues that the client is experiencing.

    Captured HTTP messages start to accumulate in the **Analysis Grid** viewer.
11. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.
12. Ensure that the **Field Chooser Tool Window** is displayed. If it is not, then select it from the **Windows** submenu that appears in the global Message Analyzer **Tools** window.
13. In the **Field Chooser**, scroll down to the **HTTP** message hierarchy, click its expansion node, and then expand the **Response** node.
14. Under the **Response** node, double-click the **StatusCode** and **ReasonPhrase** fields to add them as new columns in the **Analysis Grid** viewer column layout.

    The **StatusCode** and **ReasonPhrase** fields are populated with data values from the trace.
15. Examine the **StatusCode** and **ReasonPhrase** values of the captured data to determine possible causes of connection issues.

More Information
To learn more about common HTTP **StatusCode** and **ReasonPhrase** definitions, see the status code table in the Addendum 2: HTTP Status Codes section of this documentation.

Filtering Live Trace Session Results

This section provides examples of possible ways to filter displayed data that was collected from live traces that utilized the **Local Network Interfaces**, **Loopback and Unencrypted IPSEC**, and **Pre-Encryption for HTTPS Trace Scenarios**. The types of filters that are applied consist of Address, Port, HTTP, and TCP view **Filters**.

Applying a view Filter to Local Network Interfaces Trace Results
The hypothetical high-level issue in this example is that a Network Administrator is trying to discover the signature of virus traffic that has been infecting client computers. The administrator takes a **Local Network Interfaces** trace for an adequate time period to ensure enough data is collected on a client computer that is potentially being infected. When viewing trace results, the administrator realizes that it contains significant lower-layer noise that should be filtered out. For example, there is a lot of traffic on the client coming from a busy SQL server that is of no interest to the administrator, so he or she decides to remove it with a view **Filter** based on the SQL server IP address. There is also considerable traffic coming from the Remote Desktop Protocol (RDP) that is servicing the client's connection to a remote computer, which the administrator decides to remove with a view **Filter** based on a TCP port. In this example, the administrator can rule out SQL server and RDP

messages when searching for virus signatures because these services are normally considered cleaned by antivirus software.

To apply a view Filter to Local Network Interfaces trace results that removes unwanted traffic

1. On the client computer, perform steps 1 through 4 of the previous procedure: To filter data on a specific network interface in a Local Network Interfaces trace.

   Depending on your operating system, either the **Microsoft-PEF-NDIS-PacketCapture** provider or the **Microsoft-Windows-NDIS-PacketCapture** provider and its **Id** are displayed in the **ETW Providers** list on the **Live Trace** tab of the **New Session** dialog.

2. Verify that the **Analysis Grid** viewer is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.

3. Click the **Start** button in the **New Session** dialog to begin capturing data.

   The captured messages start to accumulate in the **Analysis Grid** viewer.

4. While Message Analyzer is capturing data, attempt to reproduce any conditions that you suspect might be causing client vulnerability to virus infection, if possible.

5. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.

6. In the text box of the default Filter panel on the Filtering toolbar, , enter the following Filter Expression to remove unwanted SQL and RDP traffic, while substituting the actual SQL server IP address for the italic value:

   ```
   *Address != 192.168.1.1 && TCP.Port != 3389
   ```

7. In the default Filter panel on the Filtering toolbar, click the **Apply** button to remove all bidirectional SQL and RDP traffic.

8. Examine the remaining data in the **Analysis Grid** viewer by performing established procedures for isolating and detecting the algorithms or behaviors that expose unique virus signatures.

Applying an HTTP view Filter to Loopback and Unencrypted IPSEC Trace Results

The hypothetical high-level issue in this example is that a Network Administrator needs to ensure that he or she can isolate *all* HTTP messages interchanged between a client computer and a web server. In this example, the administrator runs the **Loopback and Unencrypted IPSEC Trace Scenario** with the **Microsoft-PEF-WFP-MessageProvider** to ensure that all HTTP data is collected with a minimum of lower level noise, and then applies an HTTP view **Filter** to the trace results. To address this issue, the administrator uses the **Loopback and Unencrypted IPSEC Trace Scenario** to ensure capture of all HTTP traffic, rather than the **Pre-Encryption for HTTPS Trace Scenario** with the **Microsoft-PEF-WebProxy** provider, which captures browser traffic only. Therefore, the administrator can be confident that he or she has isolated all the HTTP traffic necessary to analyze whatever issues are occurring. In addition, Message Analyzer provides better functionality in these circumstances than Network Monitor, which returns only the HTTP headers rather than all message fragments as Message Analyzer does.

To apply an HTTP view Filter to Loopback and Unencrypted IPSEC trace results and examine all HTTP-related messages

1. On the server computer, perform steps 1 through 3 of the previous procedure: To filter data on a specific network interface in a Local Network Interfaces trace.

2. From the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, select the **Loopback and Unencrypted IPSEC Trace Scenario**.

   The **ETW Providers** list is populated with **Microsoft-PEF-WFP-MessageProvider** information.

3. Verify that the **Analysis Grid** viewer is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.

4. Click the **Start** button in the **New Session** dialog to begin capturing data.

   The captured messages start to accumulate in the **Analysis Grid** View.

5. While Message Analyzer is capturing data, attempt to reproduce any conditions that cause the client or web server to experience the indicated issues.

6. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.

7. In the text box of the default Filter panel on the Filtering toolbar, , enter the following atomic Filter Expression:

   ```
   HTTP
   ```

8. Click the **Apply** button on the default Filter panel to isolate all HTTP Operations and other messages that contain HTTP in the message stack.

> **NOTE**
> Alternatively, select **HTTP** from the **Viewpoints** drop-down list on the Filtering toolbar to isolate all HTTP messages at top-level, so you can see the trace results from the perspective of the HTTP protocol. In addition, you can click the **Flat Message List** button the Filtering toolbar to break out all HTTP Operations (request and response messages) into their original chronological order, similar to the Network Monitor view, for a different analysis perspective.

9. In the **Analysis Grid** viewer, click the message expansion nodes to expose the HTTP message origins tree for HTTP messages of interest.

   Alternatively, add **StatusCode** and **ReasonPhrase** columns to the **Analysis Grid** viewer from the **Field Chooser Tool Window** and execute the **Group** command on the **StatusCode** and **ReasonPhrase** columns to create groups of identical HTTP status codes, for ease of analysis.

10. Examine the filtered data to determine possible causes of any issues experienced by the client or web server, as appropriate. See Addendum 2: HTTP Status Codes for HTTP status code information.

Applying TCP view Filters to Loopback and Unencrypted IPSEC Trace Results

The hypothetical high-level issue in this example is that a Network Administrator has one or more client computers that are believed to be experiencing TCP connection and/or data transmission problems such as the following:

- Dropped packets or lost TCP segments.
- Slow data transmission rates related to TCP window size and scaling.
- Incomplete connection request handshakes.

In this scenario, the administrator identifies common TCP connection and data transmission issues by writing and applying various Filter Expressions to expose data that can reveal the potential causes of such problems. The procedure that follows provides examples of Filter Expressions that the administrator applies to isolate the values of various TCP fields, to determine whether any problems exist in these areas.

> **TIP**
> When troubleshooting TCP connection issues, Message Analyzer users should learn the meaning of all the TCP analysis flags so they can write filters to isolate and analyze the information they provide.

To apply TCP view Filters to Loopback and Unencrypted IPSEC trace results and expose TCP diagnostics

1. On the server computer, perform steps 1 through 3 of the previous procedure: To filter data on a specific

network interface in a Local Network Interfaces trace.

2. From the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, select the **Loopback and Unencrypted IPSEC Trace Scenario**.

   The **ETW Providers** list is populated with **Microsoft-PEF-WFP-MessageProvider** information.

3. Verify that the **Analysis Grid** viewer is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.

4. Click the **Start** button in the **New Session** dialog to begin capturing data.

   The captured messages start to accumulate in the **Analysis Grid** View.

5. While Message Analyzer is capturing data, attempt to reproduce the conditions where TCP connection issues are occurring on target client computers.

6. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.

7. If the **Field Chooser Tool Window** is not displayed, click the **Add Columns** button in the **Analysis Grid** viewer toolbar to display it.

8. In the **Field Chooser** window, scroll to the **TCP** protocol and click its expansion node to expose the TCP message hierarchy.

9. In the **TCP** message hierarchy, expand the **Segment** node and then double-click the **TCPDiagnosis**, **SYN** flag, **ACK** flag, **Window**, and **Options** fields to add them as new data columns in the **Analysis Grid** viewer column layout.

> **NOTE**
> Adding these fields to the column layout enables you to view the corresponding values of these fields as you apply the filters throughout this procedure.

10. In the text box of the default Filter panel on the Filtering toolbar, enter the following Filter Expression:

    ```
    TCP#DiagnosisLevels
    ```

11. On the default Filter panel , click the **Apply** button to filter the **Loopback and Unencrypted IPSEC** trace results and verify whether you have any TCP-related errors.

12. If the filter in the previous step exposes TCP errors in the **TCPDiagnosis** column of the **Analysis Grid** viewer, you can isolate and view this data in several ways, as described in the steps that follow.

13. To isolate lost or out-of-order TCP segments, enter the following Filter Expression in the text box of the default Filter panel and then click **Apply** to start the filtering process:

    ```
    *TCPDiagnosis contains "Segment-Lost"
    ```

    If you want to look at top-level TCP messages only and use a quicker performing filter, apply the following Filter Expression to the trace:

    ```
    \TCP::TCPDiagnosis contains "Segment-Lost"
    ```

14. To focus on TCP message **Window** size, which can reduce the data transmission rate when too small (see TCP Category), apply the following filter to the trace, while adjusting the italic **Window** size value in the filter as appropriate:

    ```
    TCP::WindowScaled < 1000
    ```

15. To verify whether **WindowsScaleFactor** is operative or whether other TCP **Options** are correctly set for optimal performance, apply the following filter to return *only* TCP messages that have **Options** set:

    ```
    TCP::Options ~= nothing
    ```

> **NOTE**
> If you are interested in examining the IP conversations where the TCP options were negotiated, along with the TCP ports that carried the traffic, add an IP **Network** column and a TCP **Transport** column to the **Analysis Grid** viewer column **Layout** and then perform the **Group** function on these columns to quickly reorganize the data into common groups. If you use this method, make sure to right-click the **Network** and **Transport** group labels and select the **Expand All Groups** menu command to expose all the data.

When you are viewing the TCP **Option** configuration, you might want to do the following:

- Ensure that **Selective Acknowledgement (SACK)** is enabled, to minimize TCP retransmits.
- Verify that the **Maximum Segment Size (MSS)** is set to a reasonable value to reduce fragmentation.
- Ensure that the **WindowsScaleFactor** function is operative and set to a reasonable value for the **Window** size setting, to enable automatic resize of the receive TCP **Window** as necessary.

16. To display statistics associated with all the TCP three-way handshakes in a set of trace results, execute the **TCP Three-Way Handshake** pattern expression in the **Pattern Match** viewer, which is accessible from the **New Viewer** drop-down list on the global Message Analyzer toolbar. When the **Pattern Match** viewer is open, click the **TCP Three-Way Handshake** pattern expression in the **AVAILABLE PATTERNS** list. If any matches are found, they appear in the Matched pattern selector in the **MATCHES** pane of the viewer. To review the discovered data, click the Matched pattern selector to display all instances of TCP three-way handshake patterns and the associated TCP data.

More Information
To learn more about the **TCP Three-Way Handshake** and other pattern expressions, see [Understanding Message Pattern Matching] (understanding-message-pattern-matching.md).

Otherwise, you can perform manual filtering such as the following in an attempt to locate incomplete TCP three-way handshake patterns in a trace along with surrounding messages potentially related to a failure:

- Click the **Viewpoint** drop-down list on the Filtering toolbar and then select the **TCP Viewpoint** to display all TCP messages at top-level.
- Click the **Find Message** button on the **Analysis Grid** viewer toolbar.
- Enter any of the following Filter Expressions in the **Find Message** text box and then successively click the **Find** binoculars icon to locate messages that meet the specified filtering criteria:

  TCP::TCPDiagnosis contains "Segment-Lost"
  TCP::AcknowledgementNumber==0
  TCP::Flags.Syn==true && TCP::Flags.Ack==true

You can use the method described immediately above to find TCP messages that participated in incomplete three-way handshake operations. By locating the TCP messages that contain TCP **Options** and broken or incomplete patterns of **SYN**, **ACK**, **SequenceNumber**, and **AcknowledgementNumber** field values, you may be able to determine where failures have occurred.

For reference, the following table specifies the pattern of **SYN** and **ACK** field values, with the relative **SequenceNumber** and **AcknowledgementNumber** value representations, that exposes the signature of a three-way handshake pattern that successfully opened a TCP connection:

Table 25. Three-way Handshake Signature

COMPUTER NODE	MESSAGE SENT	SYN FLAG VALUE	ACK FLAG VALUE	SEQUENCENUMBER	ACKNOWLEDGEMENTNUMBER	TCP OPTIONS
Sending	Connection request	True	False	x	0	Yes
Receiving	Request acknowledgement	True	True	y	x+1	Yes
Sending	Sync acknowledgement	False	True	x+1	y+1	No

> **NOTE**
> Keep in mind that incomplete handshakes can appear in a trace if the capture time frame did not synchronize with a TCP transmission or if the network dropped packets.

Applying a view Filter to Pre-Encryption for HTTPS Trace Results

The hypothetical high-level issue in this example is that a Network Administrator of a site has web clients that complain about slow responses when attempting to connect with and retrieve data from one or more web servers. With the use of the Message Analyzer **Pre-Encryption for HTTPS Trace Scenario** running on a client computer that is having the most issues, the administrator can quickly determine which servers are having connection or performance problems. In the procedure that follows, the administrator applies HTTP view **Filters** against data that is displaying in the **Analysis Grid** viewer, to obtain a clear picture of which servers are experiencing performance issues. To this end, the view **Filters** help the administrator accomplish the following:

- Isolate HTTP operations where server responses are slow, as indicated by high **ResponseTime** values.
- Isolate HTTP **StatusCode** values that were received, along with accompanying **ReasonPhrases**, to pinpoint any client or server errors that occurred during HTTP message exchanges.

To apply an HTTP view Filter to Pre-Encryption for HTTPS trace results and expose HTTP status codes

1. On the client computer, perform steps 1 through 3 of the previous procedure: To filter data on a specific network interface in a Local Network Interfaces trace.
2. From the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, select the **Pre-Encryption for HTTPS Trace Scenario**.

   The **ETW Providers** list is populated with the **Microsoft-Pef-WebProxy** provider information.
3. Verify that the **Analysis Grid** viewer is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.
4. Click the **Start** button in the **New Session** dialog to begin capturing data.

   The captured messages start to accumulate in the **Analysis Grid** View.
5. While Message Analyzer is running, perform some HTTP requests from client computers that are experiencing slow connection or data retrieval problems with one or more web servers that are suspected of being overburdened and performing poorly.

   The captured HTTP messages start to accumulate in the **Analysis Grid** View.
6. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.

7. Ensure that the **Field Chooser Tool Window** is displayed. If it is not, click the **Add Columns** button on the **Analysis Grid** viewer toolbar to display it.

8. In **Field Chooser**, scroll to the **HTTP** protocol and click its expansion node to expose the **HTTP** message hierarchy.

9. In the **HTTP** message hierarchy, locate the **StatusCode** and **ReasonPhrase** fields under the **Response** node and double-click each one to add it as a new column in the **Analysis Grid** viewer column **Layout**.

10. In the **Field Chooser** window, open the **Global Annotations** node and then double-click the **ResponseTime** annotation to add it as a new column in the **Analysis Grid** viewer.

11. Click the **Viewpoint** drop-down list on the Filtering toolbar and select **HTTP** to apply an HTTP viewpoint that pushes all HTTP messages to top-level with no messages above them, for ease of analysis.

12. In the text box of the default Filter panel on the Filtering toolbar, enter the following Filter Expression:
    `#ResponseTime >=1`

13. On the default Filter panel, click the **Apply** button to apply the filtering and isolate the HTTP operations where the server response to client requests is greater than or equal to 1 second.

    Operations that meet the **ResponseTime** criteria of the applied filter would likely indicate excessive server response times, providing that the **TimeElapsed** for the entire operation is a reasonable value. Note that if the **TimeElapsed** is also taking a while, this could be an indication of network issues rather than a slowly responding server.

> **NOTE**
>
> The **ResponseTime** is the difference between the **Timestamp** value of the first response message from the server minus the last **Timestamped** value in the request message stack.

14. Click the **Show Column Filter Row** icon in the upper-left corner of the **Analysis Grid** viewer to display the amber-colored **Column Filter** text box row.

15. Isolate specific HTTP **StatusCodes** in the trace by typing “10”, “20”, “30”, “40”, or “50” into the **Column Filter** text box beneath the **StatusCode** column header, as appropriate, or enter a specific **StatusCode** for which you are looking, for example “408”.

    The column is filtered according to the value you specified.

16. Examine the different **StatusCodes** and **ReasonPhrases** to determine any problem areas that might provide some additional indications as to why the web server is performing poorly. To review the definitions of key **StatusCodes** and **ReasonPhrases**, refer to the Addendum 2: HTTP Status Codes section of this documentation.

17. Analyze the **TimeElapsed** column values in the **Analysis Grid** viewer to verify whether overall HTTP operations are taking a long time.

    For example, you could type “1.”, “2.”, “3.”, and so on, into the **Column Filter** text box beneath the **TimeElapsed** column to filter for messages that have an elapsed time of 1 second or greater.

Filtering with Color Rules and Exposing Diagnostics for TCP and SMB

The hypothetical high-level issue in this example is that a Network Administrator needs to be aware of faulty or erratic client connections or file request failures that occur during file server access, by highlighting errors representing common TCP connection issues and SMB status. In the example that follows, the administrator runs the **Network Tunnel Traffic and Unencrypted IPSEC Trace Scenario** to take advantage of the

Microsoft-PEF-WFP-MessageProvider that captures data above the IP/Network Layer, along with a **Session Filter** that isolates traffic on SMB port 445. The administrator then does the following:

- Creates a new right-gradient **Color Rule** to highlight TCP diagnostic information.
- Adds a **TCP.TCPDiagnosis** column to enable quick examination of TCP errors that appear in the trace.
- Creates a new left-gradient **Color Rule** to highlight SMB error status.
- Adds an **SMB2.ErrorResponse.Header.Status.Value** column to enable quick examination of SMB2 error values that appear in the trace.
- Uses the **Color Rule** highlighting as a prompt to perform various operations for error analysis.

Applying Color Rules to Network Tunnel Traffic and Unencrypted IPSEC Trace Results and Exposing Diagnostics

1. On the client computer, perform steps 1 through 3 of the previous procedure: To filter data on a specific network interface in a Local Network Interfaces trace.
2. From the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog, select the **Network Tunnel Traffic and Unencrypted IPSEC Trace Scenario**.

   The **ETW Providers** list is populated with **Microsoft-PEF-WFP-MessageProvider** information.
3. In the **Session Filter** text box of the **New Session** dialog, enter the following Filter Expression:

   ```
   TCP.Port == IANA.Port.SMB
   ```
4. Verify that the **Analysis Grid** viewer is selected in the **Start With** drop-down list in the **New Session** dialog. If it is not, then select it.
5. Click the **Start** button in the **New Session** dialog to begin capturing data.
6. While Message Analyzer is capturing data, perform some file share access from the client computer to a file server or other location where access problems are known to be occurring.

   The messages captured on the SMB port start to accumulate in the **Analysis Grid** viewer.
7. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.
8. Create and save a new **Color Rule** with a right-gradient pattern and the following Filter Expression, to highlight any TCP messages where Application-type error events might have occurred, such as lost or incomplete segments, missing three-way handshakes, duplicate ACKs, retransmits, and so on, to expose any possible TCP connection issues:

   ```
   TCP#DiagnosisTypes
   ```

   If TCP diagnosis errors occurred in your trace, they are immediately flagged in the **Analysis Grid** viewer by the **Color Rule** configuration that you created. If you leave this **Color Rule** applied, these errors will continue to be flagged in the **Analysis Grid** viewer until you specifically disable the rule.

   For more information about creating and saving **Color Rules**, see Using and Managing Color Rules.
9. In the **Field Chooser Tool Window**, navigate the TCP message hierarchy to the **TCP.Segment** level and add the **TCPDiagnosis** field as a new column in the **Analysis Grid** viewer by using the right-click **Add as Column** command in the **Field Chooser** context menu, so the you can easily view TCP error descriptions.
10. Create and save a new **Color Rule** with a left-gradient pattern and the following Filter Expression, to highlight SMB **Status** errors that might have occurred during file share access:

    ```
    SMB2.ErrorResponse.Header.Status.Value > 0
    ```

    This **Color Rule** is also immediately applied to your trace results, in the previously indicated manner.

> **NOTE**
>
> You can leave these **Color Rules** in the applied state so that any messages with TCP diagnostics and SMB errors are automatically displayed whenever you run a trace. Unless you specifically disable **Color Rules**, they remain enabled for all Message Analyzer **Trace Scenarios** with which you capture data.

11. In the **Field Chooser** window, navigate the SMB message hierarchy to the **SMB2.ErrorResponse.Header.Status** level and add the **Value** field as a new column to the **Analysis Grid** viewer, so you can examine the values of any SMB errors that might have occurred.

12. Perform the following operations to expose TCP diagnosis messages for analysis purposes:

    - Display TCP messages at top-level to improve analysis by applying the **TCP Layer Viewpoint** from the **Viewpoints** drop-down list on the Filtering toolbar.

    - Analyze TCP error messages as appropriate to determine various TCP issues.

      For example, apply a **Column Filter** to the **TCPDiagnosis** column in the **Analysis Grid** viewer with the text "retransmit", to isolate all error messages that contain the text "Retransmitted" in the description. Then apply an "ack" **Column Filter** to isolate all error messages that contain the text "Dup-Ack" in the description. You can do the same for "Segment-Lost".

13. Perform the following operations to expose SMB error messages and other statistics for analysis purposes:

    - Break apart the SMB2 request and response messages from their Operation nodes by applying the **SMB/SMB2 Disable Operations Viewpoint** from the **Viewpoint** drop-down list on the Filtering toolbar.

    - Focus on SMB2 error messages by applying the filter `SMB2.ErrorResponse.Header.Status.Value > 0` as a view **Filter** that you specify in the text box of the default Filter panel.

      To view the filenames where such errors occurred, add an SMB2 **QueryInfoRequest.Filename** column to the **Analysis Grid** viewer with the use of the **Add as Column** context menu command in **Field Chooser**. You can then correlate the errors and filenames.

      At this point, you might also invoke one of the SMB **Layouts** for the **Chart** viewer to obtain further statistics for analysis. For example, you might select the **SMB File Stats** view **Layout** from the **Chart** drop-down list in the **New Viewer** menu on the global Message Analyzer toolbar.

Saving Message Data

2 minutes to read

After you have loaded data through a Data Retrieval Session or captured data in a Live Trace Session and displayed it in a viewer such as the **Analysis Grid** viewer, in most cases you will manipulate the data in some way to isolate specific information and analyze the issues on which you are working. For example, you might have modified your Data Retrieval Session or Live Trace Session results data in one or more of the following ways:

- Displayed additional message fields by adding specific columns to the **Analysis Grid** viewer to contain new field data.
- Applied a view **Filter** to isolate specific message data.
- Created **Color Rules** for a data set, to highlight messages based on various protocol names, field values, states, and so on.
- Set **Viewpoints** or changed a viewer **Layout**.

However, although you may have modified your data display in an Analysis Session, Message Analyzer will allow you to save only certain settings that you configured for a set of trace results. These settings include the following:

- **Filtered data** — you can save the results of applying a view **Filter** to a set of messages. You can also save the results that display in a separate **Analysis Grid** viewer tab, for example, a set of messages that display in an **Analysis Grid** viewer tab after you have double-clicked a visualizer component in another data viewer such as the **Protocol Dashboard** or some **Chart** viewer **Layout**.
- **Selected messages** — you can save any set of messages that are currently highlighted.
- **Bookmarks** — you can save bookmarks that you configured in a set of trace results, but only if you save your data in the Message Analyzer native .matp format.
- **Comments** — you can save comments that you configured in a set of trace results, but only if you save your data in the Message Analyzer native .matp format.
- **Time Shifts** — you can save a **Time Shift** that you applied in a set of trace results, but only if you save your data in the Message Analyzer native .matp format.

> **NOTE**
> In a future release of Message Analyzer, you may be able to save more settings.

What You Will Learn

In this section, you will learn about saving session data to the default location, the selection options for saving messages, specifying file formats, and naming files, as indicated by the topics below.

In This Section

Saving Session Data — learn how to use the **Save As/Export Session** dialog to save trace results data in the *.matp file format, or to export trace results to a *.cap file for compatibility with other applications, such as Network Monitor.

Selecting Messages to Save — review the save options that are available, which includes saving **All Messages**, **Filtered Messages**, and **Selected Messages**.

Saving Files in Native Format — learn about saving data in the default native file formats, saving settings such as comments and bookmarks, and the advantages of reloading data from trace files that you save in the default *.matp format. Also learn about external compatibility with other applications, such as Wireshark.

Naming Saved Files — review some tips and suggestion on naming the trace and log files that you save.

Saving Session Data

5 minutes to read

Message Analyzer enables you to save message data that is loaded from any Data Retrieval Session or captured in any Live Trace Session. If you start a new Data Retrieval Session and load data from one or more saved message files, or you capture messages in a Live Trace Session, you can save certain modifications that you made to the data set if you save in the Message Analyzer native parsed (.matp) file or .cap file formats only. The modifications to a set of trace *results* that can be saved are described in Saving Message Data. Note that the capture file (.cap) format can be viewed in other applications such as Microsoft Network Monitor.

If you open a supported message file (see Locating Supported Input Data File Types) of non-native format, for example with the **Open** feature, and you manipulate the data, you can save the resulting data set in one of the previously indicated file formats only. Also, if you load data from a saved file in parsed .matp format or any other supported format, and you need to resave it after manipulating the data, you can only save it again in the same .matp format, or you can export it as a .cap file. If you start a new Live Trace Session, you can save the session data that you capture in either of the indicated file formats.

Using the Save As/Export Session Dialog

When you are ready to save your data, click the **Save As** item in the Message Analyzer **File** menu to display the **Save/Export Session** dialog that is shown in the figure that follows. This dialog provides you with the option to save all messages, a message set that resulted from applying one or more filters or other operations, or a selected set of messages from the Data Retrieval Session or Live Trace Session results.

> **NOTE**
>
> Message Analyzer also opens the **Save/Export Session** dialog when you click the **Save Trace (Ctrl+S)** icon in the upper left corner of the Message Analyzer user interface.

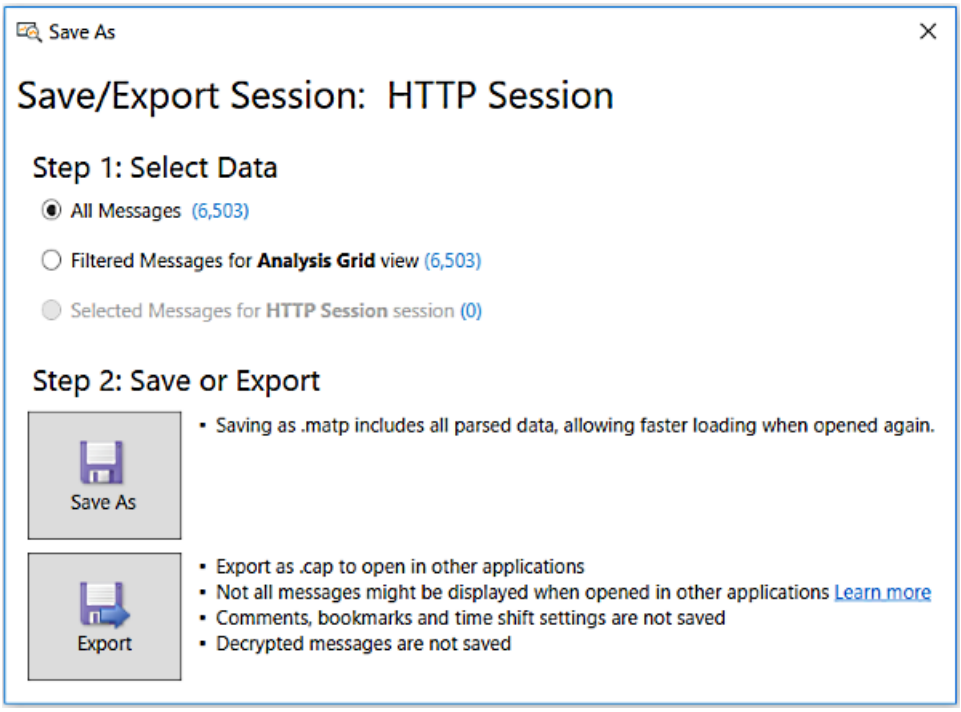

Figure 68: Save/Export Session dialog

For example, the session results that you save might be one of the following:

- Data from multiple log files imported from a Data Retrieval Session that you consolidated and processed for analysis purposes.
- Analyzed data from a Data Retrieval Session or Live Trace Session, where you manipulated the data results, for example, by applying a view **Filter** and a **Time Shift**.
- Loaded trace data to which you applied a **Session Filter** or **Time Filter**, to narrow down the focus to a specific set of messages or a window of time, respectively.
- Live trace data to which you applied a **Session Filter** or **Fast Filter**, to reduce trace results to a focused set of messages with specific properties.
- Data captured in scenarios that used ETW system providers, keyword and level filters, adapter filters, or other special filtering and modifiers.

> **TIP**
> When you want to save the results of a Live Trace Session that has never been saved before, use the **Save As** command from the Message Analyzer **File** menu to display the **Save/Export Session** dialog. If you want to save the results of changes you have made to an existing session, for example, where you have applied a **Time Shift** or added **Bookmarks** and **Comments**, use the **Save** command from the Message Analyzer **File** menu to silently preserve the changes in the session file.

Saving Session Results

After loading, capturing, manipulating, and analyzing data, you can save your session results set to a single native file only for which you specify a new name, or you can overwrite an original file with the new results set. The save process aggregates all message data in the selected session (or a filtered subset of the data) into this single file no matter how many files you originally included in the input file configuration. Moreover, if you attempt to reload data from this newly saved file, Message Analyzer will display the name of the new file in the files list and retrieve all the aggregated data contained in that file after you start the Data Retrieval Session.

Message Analyzer also enables you to save session results after working with either of the following:

- Any file that you open from the **Recent File** list, which displays to the right of **Open** item when you click it in the Message Analyzer **File** menu.
- Any logs or files that you load into Message Analyzer through the **Open** feature.

Message Analyzer enables you to save data for one session at a time only. Even if you have multiple viewer tabs open for a particular session, only one data set will be saved for that session. If you save the session data and then reload it through a Data Retrieval Session, it will display in the **Analysis Grid** viewer by default, even if you had a viewer tab open for the session other than the **Analysis Grid** viewer. If you want to save data from more than one session that Message Analyzer displays, you will need to select appropriate session tabs and save the data for each session separately.

> **IMPORTANT**
> When you save data from a Live Trace Session in the native .matp file format, Message Analyzer automatically stores the OPN parsers used in the trace so that .matp files become portable. This ensures that a different Message Analyzer installation will be able to parse the messages in ported files when they are reloaded through a Data Retrieval Session or the **Open** feature.

Saving Session Data to the Default Location

After you open the **Save/Export Session** dialog and you click the **Save As** button, a **Save As** dialog opens to the following default location for saving Message Analyzer session data:

```
c:\Users\<username>\Documents\MessageAnalyzer\Traces\
```

By saving session data in this location, you are assured that you will retain all historical trace data that you previously saved, should you need to perform an upgrade or reinstallation of Message Analyzer. Note that you can change the default location in which to save your session data; however, you should not save any session data (or assets) to the following location because it is overwritten during Message Analyzer upgrades or reinstallations:

```
c:\Users\<username>\AppData\Roaming\Microsoft\MessageAnalyzer\
```

> **TIP**
> If for some reason you need to use this location for storing Message Analyzer data, you should back up any session data or assets that you have stored there, as they will be deleted during Message Analyzer installation.

More Information

To learn more about saving message data that was loaded with a **Time Filter** applied, see the procedure Apply a Time Filter to Data Loading and Save the Message Collection.

Selecting Messages to Save

2 minutes to read

Message Analyzer enables you to select specific messages from a session that you want to save by allowing you to specify one of the following save options in the **Save/Export Session** dialog:

> **NOTE**
>
> When you save data with either the **Filtered Messages** or **Selected Messages** option, the option descriptions will indicate the type of viewer or session, respectively, for which you are saving data.

- **All Messages** — enables you to save all messages that have been loaded from a Data Retrieval Session or captured in a Live Trace Session. The data that you save depends on which viewer session tab that you select.
- **Filtered Messages** — enables you to save a filtered message set that is displayed in a viewer such as the **Analysis Grid** or a **Chart** viewer **Layout**, for example, messages that display following application of a view **Filter** or **Viewpoint**. However, note that when you open a trace file that you saved, it initially opens in the **Analysis Grid** viewer, even if you saved the data while a different viewer was displaying it.
- **Selected Messages** — enables you to save only specific messages that you select in a session, for example, in the **Analysis Grid** viewer.

  Note that the message count that displays when you save with this option, as parenthetically indicated in the option description, refers to top-level parent messages only and does not specify the total count of included underlying origins messages that supported such top-level messages or operations.

> **TIP**
>
> If you want to select specific messages to save, you can do so by selecting a message in the **Analysis Grid** viewer and using the keyboard shortcut `Shift+Down Arrow` until you highlight all the messages you want include in the save. You can then right-click the message selection and select the **Save Selected Messages...** command in the context menu. Thereafter, the Windows **Save As** dialog opens to the default save location for Message Analyzer traces. You can then specify a file name in the .matp format and save your selected messages by clicking the **Save** button in the dialog.
>
> Note that you can also export any messages you want with the use of the **Export** feature on the **Analysis Grid** viewer toolbar. However, this feature saves your messages in comma separated value (CSV) format.

Saving Files in Native Format

2 minutes to read

In Message Analyzer, you can save your data in the Message Analyzer Trace Parsed (.matp) format only, or you can export your data in the Network Monitor Capture (.cap) file format. The .matp file format is native to Message Analyzer, meaning that only Message Analyzer can open these files. However, .cap files can be opened by Microsoft Message Analyzer, Microsoft Network Monitor, and other protocol analyzers. When you save a trace file with Message Analyzer, the Windows **Save As** dialog opens to the default location, which is the following:

```
C:\Users\<username>\Documents\MessageAnalyzer\Traces
```

However, note that you can change the location for saving traces and logs to any location you choose.

Saving Settings

When you save trace files, certain types of settings are saved with them. For example, in the case of .matp files, the session **Name**, **Bookmarks**, **Comments**, and **Time Shift** settings are preserved in subsequent openings of this file type. However, when you export data as a Network Monitor Capture (.cap) file, **Bookmarks**, **Comments**, and **Time Shift** settings are not persisted in subsequent openings of this file type. Also, the session **Name** is persisted only when .cap files are opened with the **Recent** or **Open** features, which are accessible from the Message Analyzer **File** menu.

> **IMPORTANT**
>
> When you save a Message Analyzer trace file in the .cap format, be aware that there may be some interoperability issues with certain applications when attempting to open a .cap file in such an application. For more information about the media types that Message Analyzer supports when saving to the .cap file format, see Compatibility with Exported CAP Files.

Reloading Data from Supported File Formats

When you save data in the parsed format (.matp) format, the files can be large in size and may consume considerable disk space. However, they are very fast to reload into Message Analyzer because the data does not have to be reparsed. Although you cannot save Message Analyzer trace data as an unparsed file, you can still load data from such files in the .matu format as described in Locating Supported Input Data File Types. The only tradeoff is the parsing time that you incur when data from these files is being loaded into Message Analyzer.

> **NOTE**
>
> If you load data from a log such as an .etl file or any other supported file type in non-native format and you need to save the results following data manipulation and analysis, you can save the data in the .matp or .cap file format only.

Compatibility with Exported CAP Files

2 minutes to read

As previously stated, Message Analyzer enables you to export trace data in the .cap file format, which is the native file format for Microsoft Network Monitor. If you export Message Analyzer data to the .cap file format, it enables Network Monitor and other protocol analysis tools that support this format to open such a file. However, Message Analyzer only supports certain media types when exporting data to the .cap file format. Therefore, if you are saving trace data that contains frames that are unsupported, that data will not be exported. In addition, if no frames of the supported type are found in the trace data being exported, Message Analyzer will not export any messages and displays the following **Trace Save Error** message:
None of the messages could be written to file, therefore the file was not created

The media types that are supported by Message Analyzer when exporting trace data in .cap file format are described in the table that follows:

Table 26. Supported Media Types for .Cap File Exports

MEDIA TYPE	VALUE
Cap frame	0
Ethernet frame	1
ETW provider message (NetEvent)	0xFFE0

See Also

Saving Files in Native Format

Naming Saved Files

2 minutes to read

When saving trace data in one of the Message Analyzer native file formats, you are advised to consider file naming conventions. The following elements can be represented in saved message data and can therefore impact how you name the file:

- The stack level at which you ran a trace or the **Trace Scenario** you used.
- The provider/s and/or provider settings you specified in the session configuration.
- The session filter or other filters you applied when capturing or loading data.
- The name you specified for the session.
- The type of data being displayed or aggregated, such as trace data or logs, respectively.
- The way in which you processed or analyzed the data, for example, the filtering you applied to your trace results, or the resulting data set configuration.
- The data viewers you employed.
- The problem you solved.

File Naming and Searchability

Considering the variability of these factors, the content of each saved message file is almost certain to be unique. Therefore, the name of any message file that you save should be representative of its unique content, to facilitate ease of recognition when you revisit the data at a future time, or when reviewers, decision makers, and other colleagues view it for the first time. You can also apply the same principles when naming new sessions.

If you have a well thought-out file naming strategy, it can simplify your work and save time. For example, if you are importing a large number of files that display in the files list of a Data Retrieval Session, you can quickly locate and select specific files containing data that you want to load by specifying characters, in the search box on the toolbar of the **Files** tab in the **New Session** dialog, that are unique to the file names or the directory in which the files are located. For instance, you could search for all files or directories containing the characters "SMB2 Query Info". The search then highlights all files and directories containing those characters and you can then select only those files from which to load data. This feature enables you to select any single file or group of files, in any combination, that contain the data you want to analyze following the data loading process.

Note that a Data Retrieval Session might contain data from multiple input files, so the file name under which you save the data should not necessarily reflect the name of any one file. Also, the file name under which you save data and the default session name of a Data Retrieval Session or Live Trace Session can be different.

Automating Tracing Functions with PowerShell

14 minutes to read

At times, it may be advantageous for you to automate certain Message Analyzer functions that enable you to do the following:

- Utilize enhancements to the manner in which you start and stop traces, for example, with various types of triggers such as a time trigger or process trigger.
- Gain control over the type of trace that you run, for example, a linear or circular trace.
- Run traces while you are focusing on other high-priority issues.

To enable these scenarios, Message Analyzer provides you with the capability to automate the capture of network messages through PowerShell scripting. Message Analyzer makes this feature available by providing PowerShell commandlets (cmdlets) that programmatically expose PEF message tracing functionality in the PowerShell scripting environment. Other supporting configuration capabilities are also provided in the PowerShell environment to facilitate some basic Trace Session configuration, including automation triggers that define how and when Trace Sessions are started and stopped.

What You Will Learn
In the topics of this section that are listed below, you will learn about the PowerShell automation features that are available for Message Analyzer.

Encapsulating Tracing Functionality — learn what types of tracing functions you an automate with PowerShell.

Using PowerShell Cmdlets — review the functions of the PowerShell cmdlets that can work with Message Analyzer.

Examining a PowerShell Script Example — look over the code for a PowerShell script example that does the following:

- Creates a trace session object that is configured for circular capture mode.
- Specifies a message provider that will capture the data.
- Adds a TCP trace filter that is applied 150 seconds after the trace session starts.
- Creates a time trigger that defines when the session will start and a keyboard key stroke that defines how the session will stop.
- Specifies a file name and path where trace data is stored when the session is stopped.

Accessing PowerShell Cmdlets and Help — learn how to obtain the latest version of PEF cmdlet Help.

> **IMPORTANT**
> Before you run any PEF PowerShell cmdlets, ensure that you update the PEF PowerShell cmdlet help for Message Analyzer, as described in Accessing PowerShell Cmdlets and Help.

Encapsulating Tracing Functionality

The pre-configured PowerShell commandlets (cmdlets) that are available for Message Analyzer enable you to do the following:

- Configure new Trace Sessions.
- Specify a message provider for your Trace Session.
- Run Trace Sessions in circular capture mode.
- Apply a **Trace Filter** to your Trace Session configuration by specifying a predefined or custom filter expression.
- Specify the following types of triggers that can start and stop a Message Analyzer Trace Session, or perform other functions:
  - DateTime
  - KeyDown
  - Message
  - Process
  - Event Log
  - Win32 Events
  - TimeSpan
- Receive notification when a particular condition is met, for example, when an event triggers Message Analyzer to start a Trace Session or when a Trace Session stops.
- Save a message data collection in the file system.

Using PowerShell Cmdlets

The following PowerShell cmdlets automate several Message Analyzer functions so that you can streamline your network problem solving tasks, gain more control over tracing functions, and achieve better time management. The cmdlets enable you to configure, start, stop, and save data for Trace Sessions and to specify trigger events that invoke or respond to these actions, as described below.

- **Action scripts**:
  - **Invoke-PefCustomAction** — enables you to run a PowerShell script block that invokes PEF actions. You must specify the script block you want to run and a trigger for the actions it invokes. When the trigger occurs, such as a specified date-time to start or stop a Trace Session, the specified script block is invoked. The script block can contain any custom script necessary to perform custom PEF actions, for example, a script that stops a Trace Session and sends an email at that time. To learn more about script blocks, you can invoke `Get-Help about_script_blocks` at the PowerShell command line.
  - **Save-PefDataCollection** — enables you to save a collection of messages from a Trace Session. You must specify the session you want to save and the file path for the data. You can also specify a trigger that activates the save action when a Trace Session completes, so that you can save all messages currently contained in the session. You can also save a specified number of bytes so that you can analyze the data without stopping the Trace Session.
  - **Set-PefTraceFilter** — enables you to override the **Trace Filter** that you specified in a Trace Session object that you originally created with the **New-PefTraceSession** cmdlet. You can specify a string value for the **Trace Filter** and the target Trace Session as parameters of the **Set-PefTraceFilter**

cmdlet. If you use the **Set-PefTraceFilter** to specify a **Trace Filter**, it will override any filtering value that you specified with the **Filter** parameter of the **New-PefTraceSession** cmdlet. If you do not specify a trigger for the override action, the Trace Filter that you specify will take effect immediately. However, if you do specify a trigger, you can control the point in time at which the Trace Filter is applied in the Trace Session. For example, you can set the **Filter** parameter of the **New-PefTraceSession** cmdlet to a specific value and then use the **New-PefTimeSpanTrigger** cmdlet to specify a time span after which a **Trace Filter** configured by the **Set-PefTraceFilter** cmdlet is inserted into the PEF Runtime component parsing and filtering processes.

When you start a Trace Session with the **Start-PefTraceSession** cmdlet, the **Trace Filter** that you specify with the **Set-PefTraceFilter** cmdlet functions the same way as any other **Trace Filter** configured in the Message Analyzer UI. The **Set-PefTraceFilter** cmdlet also returns the target session to enable pipelining.

> **NOTE**
> If you specify a file-based data source (such as a log file) as the message provider when creating a Trace Session with the **New-PefTraceSession** cmdlet, any filter that you specify with the **Set-PefTraceFilter** cmdlet will act as a **Trace Filter**.

- **Start-PefTraceSession** — enables you to start a Message Analyzer Trace Session and to specify a trigger for the startup action. **Start-PefTraceSession** acts as an entry point for message processing. If you do not specify a trigger, **Start-PefTraceSession** initiates a processing loop where no other PowerShell cmdlets or functions are executed until the loop ends. If the **Start-PefTraceSession** cmdlet has a trigger, it will start a message processing loop only when that trigger is fired. When a message processing loop terminates, all active Trace Sessions are stopped. You can stop a Trace Session by invoking the **Stop-PefTraceSession** cmdlet, which causes **Start-PefTraceSession** to exit the processing loop. This cmdlet also returns the target session to enable pipelining.
- **Stop-PefTraceSession** — provides the means to define how you will stop a specified Trace Session. When the session is stopped, it is terminated and the PEF Runtime state is cleaned up. You can also use this cmdlet to define the trigger action that stops a specified Trace Session, which you configure prior to starting the Trace Session. To store the data retrieved in the Trace Session, you can specify values for the `SaveOnStop` parameter when creating a Trace Session with the **New-PefTraceSession** cmdlet, or you can use the **Save-PefDataCollection** cmdlet to specify where to store retrieved data.

> **NOTE**
> When you write a PowerShell script, you typically specify the **Stop_PefTraceSession** cmdlet before the **Start-PefTraceSession** cmdlet, because the Trace Session will start as soon as you hit return at the PowerShell command line after specifying the **Start-PefTraceSession** cmdlet.

- **Trigger scripts**:
  - **New-PefDateTimeTrigger** — enables you to create a date-time trigger that you can use to start a Trace Session, stop a Trace Session, or inject a **Trace Filter** into the Trace Session at a specific time. When you associate a date-time trigger with a PEF action, the computer where the Trace Session will run sets a timer that triggers the specified PEF action when the trigger is activated.
  - **New-PefKeyDownTrigger** — enables you to create a trigger action based on keyboard input, by pressing a key that you specify with the `-Key` parameter. Note that the `Ctl+C` keyboard combination is no longer supported as a trigger for this cmdlet. You can use a specified keystroke trigger to start or stop a Trace session. When you associate this trigger with a PEF action, the PEF action occurs when the trigger fires on the computer where the Trace Session is running.

- **New-PefMessageTrigger** — provides the means to create a message trigger that you can use to start, stop, save, or filter a PEF Trace Session, for example, based on a captured message type. When you associate this trigger with a PEF action, the PEF action occurs when the trigger fires on the computer where the Trace Session is running.
- **New-PefProcessTrigger** — enables you to create a process trigger that starts a Trace Session when a process exits. For example, you might start a Trace Session after a started executable process has finished running. When you associate this trigger with a PEF action, the PEF action occurs when the trigger fires on the computer where the Trace Session is running.
- **New-PefTimeSpanTrigger** — enables you to create a timer trigger that fires after a specified time span. You can use this timer trigger to start, stop, or add a **Trace Filter** to a Trace Session when a specified interval of time elapses.
- **New-PefEventLogTrigger** — enables you to create a trigger that fires when an entry is created in the Windows Event Log.
- **New-PefWin32EventTrigger** — enables you to create a trigger that fires when a Win32 Event object is set.

- **Miscellaneous scripts**:
  - **Add-PefMessageProvider** — enables you to add one or more message providers to a specified Trace Session object that you create with the **New-PefTraceSession** cmdlet. A message provider can be a PEF message provider such as the **Microsoft-Pef-WFP-MessageProvider**, a system ETW provider, or even a file-based message source such as a log file.
  - **Add-PefMessageSource** — enables you to add various message sources to a PEF Trace Session. You can specify any of the following types of message sources for this cmdlet:
    - A saved file as a message source in a Data Retrieval Session, for example, a .cap or .matp file.
    - A log file as a message source in a Data Retrieval Session, for example, an event log (.etl) or a text log (.log) file.

    > **NOTE**
    > When specifying a text log as input to a Data Retrieval Session, you can also specify a text log configuration file for parsing the log.

  - Manifest-based PEF message providers such as the **Microsoft-PEF-WFP-MessageProvider** or **Microsoft-PEF-NDIS-PacketCapture** provider, as a message source in a Live Trace Session.
  - Manifest-based system ETW providers such as the **Microsoft-Windows-Dhcp-Client** as a message source in a Live Trace Session.
  - The object created by the **Add-PefProviderConfig** cmdlet can also serve as input to the **Add-PefMessageSource** cmdlet.
- **Add-PefProviderConfig** — enables you to add a provider to the configuration of a Live Trace Session that targets a remote host, by specifying the friendly provider name (not a GUID).

  The provider configuration is accessible by using the object that this cmdlet creates, from where you can configure provider error levels, event keywords, filters, and other provider-specific options for providers such as the **Microsoft-Windows-NDIS-PacketCapture**, **Microsoft-Pef-WebProxy**, and **Microsoft-Pef-WFP-MessageProvider**.
- **New-PefTraceSession** — enables you to create a Trace Session object that captures live data or retrieves

stored messages, for example, from a log file. You can specify whether to capture data in circular or linear mode to control how much data is held in the Trace Session. You can also configure a **Trace Filter** to focus the data retrieval action on messages that meet specific filtering criteria. If you want to save the data to a file after the Trace Session is stopped, you can do so by specifying the **SaveOnStop** parameter. For each Trace Session that you configure, you must add the message provider you want to use, such as the **Microsoft-Pef-NDIS-PacketCapture** or **Microsoft-Pef-WFP-MessageProvider**, by specifying it with the **Add-PefMessageProvider** cmdlet. To start and stop the Trace Session, you can use the **Start-PefTraceSession** and **Stop-PefTraceSession**, respectively, along with configuring any triggers that facilitate such actions.

- **New-PefTargetHost** — enables you to create a target host object that you specify as a target computer for remote tracing in a Live Trace Session. You can target and add multiple computers to a Live Trace Session with this cmdlet. Use the object that this cmdlet creates as input to the **Add-PefProviderConfig** cmdlet. Note that this cmdlet uses the current user credentials by default, although you can provide other credentials by specifying the -Credentials parameter.

  > **NOTE**
  > This cmdlet uses the Microsoft-PEF-WFP-MessageProvider, which is now enabled for remote tracing. You can use this provider to capture remote traffic, but on a remote Windows 10 host only, and while running on a Windows 8.1, Windows Server 2012 R2, or Windows 10 computer only.

  > **TIP**
  > On the Windows 10 client operating system, you can capture traffic locally or remotely in promiscuous mode (p-mode) by using the **Add-NetEventNetworkAdapter** and the **Add-NetEventPacketCaptureProvider** PowerShell cmdlets. With the **Add-NetEventNetworkAdapter** cmdlet, you can specify the –PromiscuousMode parameter for a supporting network adapter that you are adding as a filter on a remote packet capture provider. With the **Add-NetEventPacketCaptureProvider** cmdlet, you can specify the **Windows-NDIS-PacketCapture** provider to capture remote traffic and save a *.etl file locally on the remote computer. Note that you can import this file into Message Analyzer from the **Files** tab of the **New Session** dialog to retrieve the data for analysis. In a future Message Analyzer release, you may have the option to capture in the promiscuous mode directly from the user interface.

> **To learn more**, see the Add-NetEventNetworkAdapter and Add-NetEventPacketCaptureProvider cmdlets on TechNet.

Examining a PowerShell Script Example

To automate Message Analyzer network trace functionality with PowerShell, you will need to string together PowerShell cmdlets to achieve a desired result. The base cmdlet upon which all other cmdlet functionality depends is the **New-PefTraceSession** cmdlet. For example, you must use this cmdlet first to create a trace session object and then use other cmdlets to include additional configurations, such as adding the provider to use in the Trace Session, specifying an override filter, adding data saving functions, and configuring trigger actions.

The following example uses the `New-PefTraceSession` cmdlet to create a Trace Session object that is stored in the variable `$TraceSession01` and is configured for the circular capture mode. The script then uses the `Add-PefMessageProvider` cmdlet to specify the provider that Message Analyzer should use to capture data and associates the provider specification with `$TraceSession01`. Next, the `Set-PefTraceFilter` configures a "TCP" **Trace Filter** that will be applied to the trace 150 seconds after the Trace Session starts, as specified by the variable `$Trigger01`, which is configured with the `New-PefTimeSpanTrigger` cmdlet. The script then specifies two more triggers; `$Trigger02`, which configures the time at which the Trace Session will start, and `$Trigger03`, which specifies the PEF action that stops the Trace Session, which in this case is a keyboard key that is specified by the

`-Key` parameter. These triggers are then associated with the `Stop-PefTraceSession` and `Start-PefTraceSession` cmdlets, respectively. Lastly, the `Save-PefDataCollection` cmdlet specifies the trace file type (.matu) and the file name and full path where the Trace Session data will be stored at the time `$Trigger03` occurs. The `Force` parameter in this cmdlet causes the data of any existing file of the same name to be overwritten. The `Start-PefTraceSession` cmdlet then begins the session when `$Trigger02` fires.

The syntax for this functionality is specified as follows:

```
$TraceSession01 = New-PefTraceSession -Mode Circular
Add-PefMessageProvider -PEFSession $TraceSession01 –Provider "Microsoft-PEF-WFP-MessageProvider"
$Trigger01 = New-PefTimeSpanTrigger -TimeSpan (New-TimeSpan -Seconds 150)
Set-PefTraceFilter -PEFSession $TraceSession01 –Filter "TCP" -Trigger $Trigger01
$Trigger02 = New-PefDateTimeTrigger -DateTime "9/30/2013 7:00 AM"
$Trigger03 = New-PefKeyDownTrigger –Key S
Stop-PefTraceSession -PEFSession $TraceSession01 -Trigger $Trigger03
Save-PefDataCollection -PEFSession $TraceSession01 -Path <"fullTracePath\myTrace.matu"> -Force -Trigger
$Trigger03
Start-PefTraceSession -PEFSession $TraceSession01 –Trigger $Trigger02
```

More Information

To learn more about writing **Trace Filters** and other Filter Expressions, see the Writing Filter Expressions topic in this Operating Guide. Note that a **Trace Filter** in a PowerShell script performs a function that is identical to a **Session Filter** in the Message Analyzer user interface.

Accessing PowerShell Cmdlets and Help

To take advantage of the functionality provided in the previously described Using PowerShell Cmdlets for Message Analyzer, you must have PowerShell v3.0 installed. PowerShell v3.0 installs automatically with Windows 8 and later operating systems; however, if you are running Windows 7, you will need to install the Windows Management Framework 3.0 to obtain a PowerShell v3.0 installation. After you have a PowerShell v3.0 installation in place on your Windows 7 machine, you will need to run the following command to import the PEF PowerShell module into your PowerShell session:

`Import-Module PEF`

Then, to update the help, run the following command to download the latest cmdlet Help content from TechNet:

`Update-Help -Module PEF -Force –Verbose`

> **NOTE**
>
> PowerShell cmdlet help documentation is available at the PowerShell command line and also in the TechNet Library on the Message Analyzer Cmdlets site. For complete command-line syntax, parameter specifications, and usage examples, see these locations. If you want to view help at the PowerShell command line for a particular cmdlet, specify the following command string:
>
> `get-help <cmdletname>`

Managing Message Analyzer Assets

2 minutes to read

Message Analyzer provides an infrastructure that enables you to download, automatically update, and manage various types of assets that you commonly use in Message Analyzer operations. Some of the assets that you can use include view **Filters**, **Trace Scenarios**, **Analysis Grid** viewer **Layouts**, **Color Rules**, **Chart** viewer **Layouts**, **Viewpoints**, **OPN Parser** packages, and so on. With exception of the **OPN Parser** packages, these assets exist in user Libraries that are accessible from various locations in the Message Analyzer user interface. You can manage these assets from the **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu. In the **Asset Manager** dialog, you will find asset collections such as **Message Analyzer Filters**, **Message Analyzer Chart View Layouts**, and others that are included in every Message Analyzer installation. All asset collections that appear in the **Asset Manager** can be synced for periodic updates from a Microsoft web service that automatically refreshes your Libraries with the latest collection items. The infrastructure also enables you to import and export asset collections or any portion thereof, from and to a user-designated file share location, respectively. You can also create your own feeds from the **Asset Manager** dialog. The framework that contains all these capabilities is known as the Message Analyzer Sharing Infrastructure.

What You Will Learn

In the topics of this section, you will learn how to use and manage Message Analyzer Sharing Infrastructure components, which includes how to configure automatic updates for the default asset collections provided with every Message Analyzer installation, how to download asset collections, and how to manage data feeds. You will also learn about the location and structure of the user Libraries that contain your default asset collections.

In This Section

Sharing Infrastructure — learn about the Sharing Infrastructure capabilities, Message Analyzer user Library features, and the asset management features that are available for your use, which includes the Message Analyzer **Asset Manager**.

Managing Asset Collection Downloads and Updates — learn how to opt-in or opt-out of automatic updates for asset collections at first Message Analyzer startup, filter and search for asset collections, and how to download asset collections or auto-sync those collections that are not already set for automatic updates.

Managing Microsoft OPN Parser Packages — learn how to download and update various **OPN Parser** packages that Message Analyzer requires for normal operation. Note that you can manage **OPN Parser** downloads and updates with the previously mentioned status icons and synchronization options.

Managing the Default Subscriber Feed — learn about download and update processes through the default **Message Analyzer** subscriber feed, along with feed management options that include how to delete and restore this feed.

Creating Custom User Feeds — learn about the data feed configuration capabilities of the Sharing Infrastructure that enable you to configure and remove your own user data feeds.

Sharing Asset Collections on a User File Share — learn about how Message Analyzer enables you to share asset collections or any portion thereof with other users, for mutual collaboration and benefit.

Go To Procedures

To proceed directly to procedures that provide examples of managing Message Analyzer assets, see **Procedures: Using the Asset Management Features**.

Sharing Infrastructure

3 minutes to read

The Message Analyzer Sharing Infrastructure is a centralized framework that integrates specific functions of the **Asset Manager** with user **Library** asset collections in your Message Analyzer installation. Asset collections are accessible from Library drop-down lists, which contain tools that you can use to capture, filter, and manipulate the data, or change the data view, for example, with **Trace Scenario**, **Filter**, **Viewpoint**, and **Chart** viewer **Layout** Library types, respectively. In addition, a common management dialog is available so that you can manage the items in your asset collection Libraries.

Sharing Infrastructure Overview

Some of the specific entities that you can use to manage the Sharing Infrastructure consist of the following:

- **Asset Manager** dialog — access from the global Message Analyzer **Tools** menu and perform the following tasks:
  - View the list of default asset collections that are provided by Microsoft to all Message Analyzer installations.
  - Access and exercise the auto-syncing and downloading functions for the default Message Analyzer asset collections that are stored in user Libraries.
  - Create new data feeds for mutually sharing assets with others.
  - Search for various types of assets.
- **Manage** dialog — use a common dialog to manage the asset collections in various user Libraries that are integrated with the Sharing Infrastructure. Provides facilities for importing, exporting, and modifying asset collections from any Library with which you are working.
- **Add New Feed** feature — create your own custom user feed through which you can export or import asset collections for mutual sharing with others, including any items that you develop.
- **Export/Import** commands — share asset collections directly with other users by posting or retrieving them to/from a user-designated file share or other location.

Managing and Sharing User Libraries

The **Libraries** that contain your asset collections are centralized, up-dateable, user-expandable, and shareable. Managing user **Libraries** includes creating new items, organizing asset collections, and sharing them with others, as follows:

- **Manage asset collection Libraries** — you can create, reconfigure, export, import, and save asset collections to local user Libraries that are integrated with the Sharing Infrastructure, so that you, other team members, or the larger Message Analyzer community can mutually share asset collection items, including any of the default Message Analyzer items that you copy and modify. You can manage your assets from a common **Manage *<AssetType>*** dialog that displays from various **Library** drop-down lists in the Message Analyzer user interface.

> **NOTE**
>
> The **Manage <*AssetType*>** dialog functionality is common across all user Libraries; however, the actual dialog name varies depending on which Library you are managing. For example, when you are managing the **Filters** Library, the dialog is named **Manage Filter**; when you are managing the **Color Rules** Library, the dialog is named **Manage Color Rule**; and so on.

From the common **Manage <*AssetType*>** dialog, you can create new items that get added to the **My Items** category, copy and modify any of the built-in items in a Library, or delete any item in the **My Items** category. You can also select items for export configurations and import items from a file share, user feed, or other designated location. The asset collections that are provided with Message Analyzer by default are contained in *.asset files, the contents of which display in the **Libraries** that are described in User Libraries.

- **Share asset collections from a user feed** — you can configure a custom user subscriber data feed to which you can export any of the built-in Message Analyzer asset collection items or any of those that you created in the **My Items** category of a local user Library for a particular item type, so that others can import them for their own use. When you create your own feed, you must specify a **Feed Name** and a directory **Location** where you intend to export selected items or collections to make them accessible to other users.
- **Share asset collections from a file share** — you can share your asset collection Library items directly with others by using the **Export** command in the **Manage <*AssetType*>** dialog for the particular Library. You simply select the items or item categories that you want to export and you are prompted to navigate to a file share or other location where you want to post the items in an asset file for sharing with others. Similarly, you can use the **Import** command to access asset collections that others have shared.

The topics below provide further details about the **Asset Manager** functions, user asset collection Libraries, and how to manage these Libraries.

See Also

Asset Manager
User Libraries
Managing User Libraries

Asset Manager

3 minutes to read

Message Analyzer provides an **Asset Manager** dialog that enables you to manage your asset collections, which includes the following tasks:

- Downloading asset collections to your user Libraries.
- Auto-syncing asset collections so that you receive updates to specific Libraries automatically and silently in the background.
- Configuring feed locations to which you can post assets that you created for sharing with others, or retrieve assets from feeds that others have created.

The **Asset Manager** is accessible from the global Message Analyzer **Tools** menu. This dialog contains two tabs that enable you to perform the asset management tasks that are described below.

Downloads tab — displays a page that enables you to view, auto-sync updates, and download **OPN Parser Packages** or user **Library** asset collections, such as **Message Analyzer Chart View Layouts**, **Message Analyzer Charts**, **Message Analyzer Color Rules**, **Message Analyzer Correlations**, **Message Analyzer Grouping Viewer Layouts**, and so on, from the default **Message Analyzer** subscriber feed. The **Downloads** tab of the **Asset Manager** dialog is shown in the figure that follows.

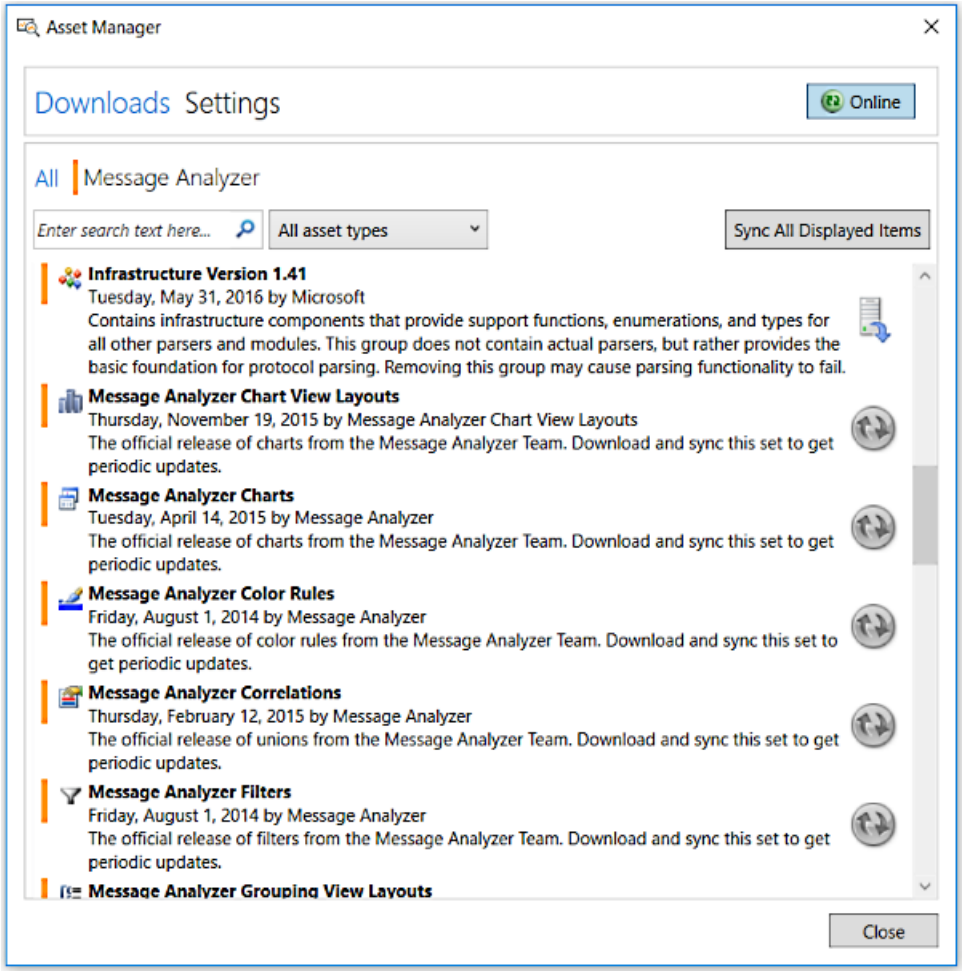

Figure 69: Asset Manager dialog Downloads tab

The features on this tab that enable you to exercise the above functions consist of the following:

- **Search box** — enables you to locate feed list items by entering search text that filters the list items.
- **All Asset Types** drop-down — enables you to select the types of assets to display in your **Message Analyzer** feed list.
- **Sync All Displayed Items** button — causes automatic update synchronization of all the default **Message Analyzer** feed list items. Thereafter, your Message Analyzer installation is synchronized with item collection and **OPN Parser** package updates from a Microsoft web service, so that you always have the latest versions. However, for this synchronization to take place, your Message Analyzer installation must be set to **Online** status, as described in the next bullet item.

> **NOTE**
>
> You have the option to synchronize **Message Analyzer** feed list items individually or you can elect to download the current configuration of a particular item and stop receiving updates. You can access these features by clicking the status icons to the right of the feed list items that are displayed on the **Downloads** page.

- **Online/Offline** button — when this button is set to **Online**, you automatically receive updates to **Message Analyzer** subscriber feed **Synced Collections** that are set to auto-sync status. When this button is set to **Offline**, you do not receive updates.

Settings tab — enables you to view a list of feeds to which you are subscribed, create new feeds for sharing item collections directly with other users, and view which asset collections and **OPN Parser** packages are set to the auto-sync state. You can also manage downloads and auto-syncing from this location. The **Settings** tab of the **Asset Manager** dialog is shown in the figure that follows.

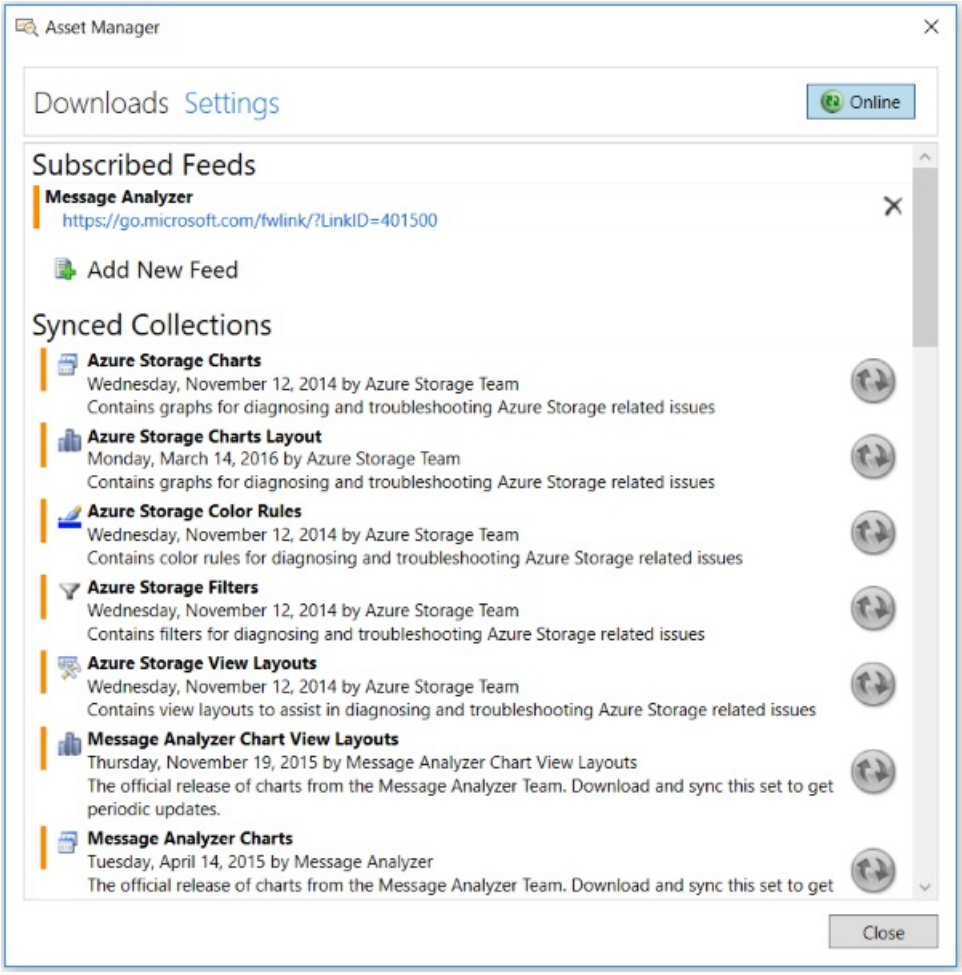

Figure 70: Asset Manager dialog Settings tab

The features on this tab consist of the following:

- **Add New Feed** — enables you to create a new feed that points to a chosen location such as a file share or web site. You would typically create a new feed for storing custom asset collections that you created so that other users may access and share them.
- **Remove Feed** — enables you to remove any feed by clicking the **X** to the right of the feed name.

> **NOTE**
> You are advised to *not* remove the default **Message Analyzer** subscriber feed.

- **Online/Offline** button — when this button is set to **Online**, you automatically receive updates to **Message Analyzer** subscriber feed **Synced Collections** that are set to auto-sync status. When this button is set to **Offline**, you do not receive updates.

See Also

Managing Asset Collection Downloads and Updates

User Libraries

8 minutes to read

Message Analyzer provides default asset collections that appear in various user Libraries in the user interface, as described in User Library Locations. From these user Libraries, you can select and apply the built-in functionality of default items to a displayed set of messages, or in the case of **Trace Scenarios**, select and apply predefined trace template functionality to capture specific data in a live trace. Moreover, you might apply a built-in view **Filter** item to trace results by selecting one from the **Library** on the default Filter panel of the Filtering toolbar. You can also expand any user Library, with the exception of the **Viewpoints** and **Parsing Level** libraries, by creating and adding your own items to it, which then appear in the **My Items** category of the associated Library. Thereafter, you can select and apply any of your own items just as you would one of the default items.

User Library Configuration

Message Analyzer enables you to **Edit**, **Delete**, set as **Favorite**, or **Create a Copy** of any item in the **My Items** category of a user Library. For items in the default **Message Analyzer** category of any user Library, you can only set the **Favorite** status or **Create a Copy** of an item.

For example, you can modify a **Color Rule**, view **Filter**, or **Pattern Expression** in the **My Items** category by right-clicking it and selecting the **Edit** command from the context menu that appears, which thereafter displays the **Edit Item** dialog that contains the configuration features you will need to modify item functionality. Note that you can also make a copy of any item in the **Message Analyzer** category of a default asset collection such as the **Color Rule**, view **Filter**, or **Pattern Expression** collection, modify it, and then save it under a new name and category, which is tantamount to creating a new item. However, although you cannot delete any of the default items from a user Library, you can delete any item in the **My Items** category of any user Library. In the case of the default asset collections such as **View Layouts**, **Trace Scenarios**, **Chart** viewer **Layouts**, and so on, modifications are limited to changing the **Name**, **Description**, and **Category** when you use the **Create a Copy** command (from the context menu that displays when you right-click a default item).

In addition, most user **Libraries** have an **Examples** subcategory under **My Items** that contains a sample that you can build upon as practice in developing your own Library items, in which you are encouraged to engage. For further details about the operations you can perform on items in each user Library, see the table in Managing User Libraries.

User Library and Sharing Infrastructure Integration

All user **Libraries** are integrated with the Message Analyzer Sharing Infrastructure and have corresponding asset collections that display in the **Asset Manager** dialog. From the **Asset Manager** dialog, you can refresh any user Library you wish by downloading an asset collection that corresponds to the particular Library, or by setting that collection to auto-sync updates, as described in the Managing Asset Collection Downloads and Updates section. You can also interact with the Sharing Infrastructure from the **Manage <*AssetType*>** dialog of any user Library only to the extent of posting and retrieving items to and from a user-configured subscriber feed, respectively. This dialog is accessible from a **Manage <*AssetType*>** menu item in each Library drop-down list, as described in Managing User Libraries, with the exception of the **Parsing Level** Library, which does not have this menu item. From the **Manage <*AssetType*>** dialog, you can **Export** any item in a Library to a file share or other location so that you can share it with others. You can also **Import** items that other users have shared to a designated location. The dialog also enables you to select which items you want to **Export** or **Import** from and to your Library and you can **Delete** items from the dialog in the **My Items** category only.

User Library Categorical Structure

By default, most asset collection **Libraries** contain the following top-level categories of items.

> **NOTE**
>
> Some Message Analyzer asset collections have no categories all, such as the **Message Analyzer Window Layouts** and **Message Analyzer Parsing Levels** collections. These assets appear as simple lists in the Message Analyzer user interface.

- **Favorites** — this category is generated when you configure any asset as a Favorite, by clicking the white star next to the asset name in the drop-down list of a particular Library. There must be at least one Favorited item for this category to exist.
- **Message Analyzer** — this category corresponds to the default **Message Analyzer** subscriber feed and contains a default asset collection for each particular user Library type. This is the category that is refreshed with periodic updates when you set corresponding asset collections to the auto-sync state from the **Asset Manager**. The **Message Analyzer** category in each Library exists by default and persists while you are auto-syncing. However, if you download an asset collection, rather than auto-sync it, all the items in the collection are moved into the **My Items** category of the Library. If you choose at some point to auto-sync the asset collection again, the **Message Analyzer** category will be restored in the Library to hold the asset collection items until such time that you perform a download once more.
- **My Items** — this category contains an **Example** subcategory that provides a sample item for development practice. However, as indicated earlier, this category is also the repository for asset collections that you download. Note that once you download an asset collection, the only way you can remove it from the **My Items** category is by using the **Manage <*AssetType*>** dialog to **Delete** it.
  - **Examples** — a **My Items** subcategory that provides a sample asset that you can work with. The sample asset is available for you to modify however you want, and is meant to help you get started with developing your own assets. After you modify the sample, you can save it to an existing category or another one that you define.

User Library Locations

Each default asset collection is included in a unique user Library in the Message Analyzer locations listed below. The formal Library names in the lists that follow are as specified in the **Asset Manager** dialog. Note that although the Library accessibility points for **Session Filters**, view **Filters**, and **Viewpoint Filters** are different, it is the same centralized Filter Expression **Library** that you access in each case:

- **Global Message Analyzer toolbar** — the following user Libraries are accessible from the global Message Analyzer toolbar:
  - **Message Analyzer Chart View Layouts** Library — **Chart** viewer **Layouts** are accessible from the **New Viewer** drop-down list on the above indicated toolbar.
  - **Message Analyzer Charts** Library — accessible from the **Charts (Deprecated)** menu in the **New Viewer** drop-down list on the specified toolbar.
  - **Message Analyzer View Layouts** Library — **Analysis Grid** viewer **Layouts** are accessible from the **New Viewer** drop-down list on the specified toolbar.
  - **Message Analyzer Window Layouts** Library — a non-categorical library that is accessible from the **Window Layout** drop-down list on the above specified toolbar.
  - **Message Analyzer Grouping View Layouts** Library — **Grouping** viewer **Layouts** are accessible from the **New Viewer** drop-down list on the specified toolbar.

- **Aliases** Library — accessible from the **Aliases** drop-down list on the specified toolbar, although this is not an asset collection that you can auto-sync or download, as currently it does not appear in the **Asset Manager** dialog.

- **Filtering toolbar** — the following user Libraries are accessible from the Message Analyzer Filtering toolbar that displays above the main analysis surface:
  - **Message Analyzer Filters** Library — accessible from the **Library** drop-down list in any Filter panel on the Filtering toolbar. Also accessible from the **Library** drop-down list on a **Viewpoint** Filter panel, but only after you apply a **Viewpoint**.
  - **Message Analyzer Viewpoints** Library — accessible by clicking the **Viewpoint** drop-down list on the Filtering toolbar.
- **Analysis Grid** toolbar — the following user Libraries are accessible from the **Analysis Grid** viewer toolbar:
  - **Message Analyzer Color Rules** Library — accessible from the **Color Rules** drop-down list on the above indicated toolbar.
  - **Message Analyzer View Layouts** Library — accessible from the **Layout** drop-down list on the above indicated toolbar.
- **Grouping** viewer toolbar — the following user Libraries are accessible from the **Grouping** viewer toolbar:
  - **Message Analyzer Grouping View Layouts** Library — accessible from the **Layout** drop-down list on the above indicated toolbar.
  - **Message Analyzer Filters** Library — accessible from the **Library** drop-down list in any Filter panel that displays on the **Grouping** viewer toolbar when you click the **Add Filter** item in the **Add Filter** drop-down list. Also accessible from the **Library** drop-down list in any **Viewpoint** Filter panel on the **Grouping** viewer toolbar, but only after applying a **Viewpoint** from the **Grouping** viewer toolbar.
  - **Message Analyzer Viewpoints** Library — accessible by clicking the **Viewpoint** drop-down list on the **Grouping** viewer toolbar.
- **Pattern Match** viewer — the **Message Analyzer Sequence (Pattern) Expressions** user Library is accessible from the **AVAILABLE PATTERNS** list in the **Pattern Match** viewer.
- **Global Message Analyzer Tools** menu — the following user Libraries are accessible from the **Tools** menu:
  - **Aliases** Library — accessible from the **Aliases** drop-down list, although this is not an asset collection that you can auto-sync or download, as currently it does not appear in the **Asset Manager** dialog.
  - **Message Analyzer Correlations** — accessible by clicking the **Unions** drop-down list in the global Message Analyzer **Tools** menu.
  - **Message Analyzer Window Layouts** Library — a non-categorical library that is accessible from the **Window Layout** drop-down list in the **Windows** submenu that appears in the global Message Analyzer **Tools** menu. However, this is not an asset collection that you can manage with the **Manage <*AssetType*>** dialog.
- **Session configuration** — the following user Libraries are accessible from the **New Session** dialog:
  - **Message Analyzer Trace Scenarios** — accessible by clicking the **Select Scenario** drop-down list on the **Live Trace** tab of the **New Session** dialog during session configuration.

- **Message Analyzer Filters** — accessible by clicking the **Library** drop-down list above the **Session Filter** text box of the **New Session** dialog during session configuration. This is the same centralized **Library** for all Filter Expressions that is also shared by the **Filter** Library in all Filter panels on the Filtering toolbar.
- **Message Analyzer Parsing Levels** — a non-categorical library that is accessible from the **Parsing Level** drop-down list in the **New Session** dialog, although this is not an asset collection that you can manage with the **Manage <*AssetType*>** dialog.

See Also

Managing User Libraries

Managing User Libraries

9 minutes to read

Message Analyzer enables you to manage any user Library by using a common and centralized set of management features. These features consist of commands and configuration settings that are available from user Library drop-down lists and from the common **Manage <*AssetType*>** dialog that is accessible from such drop-down lists and other locations.

You can use the Library drop-down lists to do the following:

- Apply the functionality of asset collection Library items to select specific information when loading data through a Data Retrieval Session or while capturing data in a Live Trace Session.
- Apply Library items to a set of trace results in an Analysis Session to manipulate your message data, so you can achieve a focused analytical perspective.
- Work with example items.
- Create new asset collection Library items.
- Execute the commands that are specified in the table of this section for items in the indicated user Libraries.
- Open the **Manage <*AssetType*>** dialog to access additional management functions, such as the **Export** and **Import** commands.

You can use the **Manage <*AssetType*>** dialog to import and export selected asset collection Library items from and to a user file share or other location, respectively, for mutually sharing Library items with others. In addition, the **Manage <*AssetType*>** dialog also enables you to perform other operations, for example, modifying and deleting existing items in certain categories.

Centrally Managing User Libraries

The central locations from where you can manage the items in any asset collection Library and the specific commands that are available from these locations are described in this section. The list that follows describes the management locations and indicates how to access commands:

- **Library drop-down list** — from the drop-down list of a Library, you can access a right-click context menu that provides different sets of commands, depending on whether you right-click an item in the default **Message Analyzer** category or an item in the **My Items** category of an asset collection Library. The available context menu commands are described ahead.
- **Manage <*AssetType*>** dialog — from this dialog, which is common to all asset collection Libraries, toolbar commands are available for importing and exporting items in any category from and to a file share or other location. Other commands that you can access from this dialog are described ahead.

An example of the management dialog is shown in the figure that follows, which in this case is the **Manage Column Layouts** dialog that provides various column **Layouts** that you can select for the **Analysis Grid** viewer. Note that the right-click context menu commands for the **File Sharing Perf SMB2/SMB** view **Layout** are displayed in the figure.

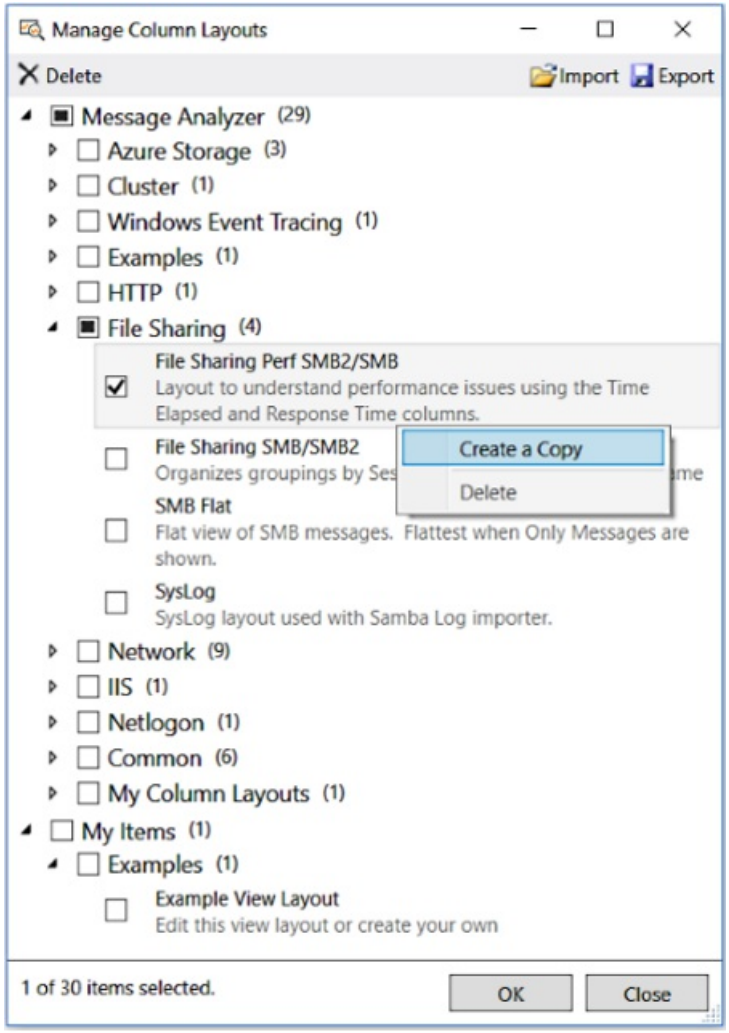

Figure 71: Manage Column Layouts dialog

The management features and context menu commands that are available in the user Library drop-down lists for Message Analyzer asset collection Libraries are described in the table that follows. Note that the available context menu commands specified in this table are also accessible from the **Manage <*AssetType*>** dialog for corresponding user Libraries, but in some cases with the exception of the **Favorites** command.

Table 27. Asset Collection Library Drop-Down Management Features

ASSET COLLECTION LIBRARY NAME	LIBRARY DROP-DOWN LIST (Y/N)	MANAGEMENT ENTITIES	EXAMPLE ITEM (Y/N)	AVAILABLE CONTEXT	MENU COMMANDS
				Message Analyzer Library category commands	**My Items** Library category commands

ASSET COLLECTION LIBRARY NAME	LIBRARY DROP-DOWN LIST (Y/N)	MANAGEMENT ENTITIES	EXAMPLE ITEM (Y/N)	AVAILABLE CONTEXT	MENU COMMANDS
Message Analyzer Filters	Yes	**New Filter** and **Manage Filters** items; **Import** and **Export** commands For more information about creating new filters, see Working with Session Filters in a Live Trace Session.	Yes	**Create a Copy**, **Delete**, **Favorite**	**Edit**, **Create a Copy**, **Delete**, **Favorite**
Message Analyzer Viewpoints	Yes	**Manage Viewpoints** item; **Import** and **Export** commands For more information about **Viewpoints**, see Applying and Managing Viewpoints.	No	None	None
Message Analyzer Color Rules	Yes	**New Color Rule** and **Manage Color Rules** items; **Import** and **Export** commands For more information about creating **Color Rules**, see Using and Managing Color Rules.	Yes	**Create a Copy**	**Edit**, **Create a Copy**, **Delete**

ASSET COLLECTION LIBRARY NAME	LIBRARY DROP-DOWN LIST (Y/N)	MANAGEMENT ENTITIES	EXAMPLE ITEM (Y/N)	AVAILABLE CONTEXT	MENU COMMANDS
Message Analyzer View Layouts	Yes	**Save Current Layout As..**, **Save Current as Default User Layout**, **Load Default User Layout**, **Restore Application Default Layout**, and **Manage Layout** items; **Import** and **Export** commands For more information about saving and restoring **Analysis Grid View Layouts**, see Applying and Managing Analysis Grid Viewer Layouts.	Yes	**Create a Copy**, **Favorite**	**Edit**, **Create a Copy**, **Favorite**, **Delete**
Message Analyzer Grouping View Layouts	Yes	**Save Current Layout As..**, **Save Current as Default User Layout**, **Load Default User Layout**, **Restore Application Default Layout**, and **Manage Layout** items; **Import** and **Export** commands For more information about saving and restoring **Grouping View Layouts**, see Managing Grouping View Layouts.	No	**Create a Copy**, **Favorite**	**Edit**, **Create a Copy**, **Favorite**, **Delete**

ASSET COLLECTION LIBRARY NAME	LIBRARY DROP-DOWN LIST (Y/N)	MANAGEMENT ENTITIES	EXAMPLE ITEM (Y/N)	AVAILABLE CONTEXT	MENU COMMANDS
Message Analyzer Pattern Expressions	Yes	**Create Pattern** and **Manage Pattern Expressions** items; **Import** and **Export** commands For more information about creating **Pattern Expressions**, see the Pattern Match Viewer topic.	Yes	**Create a Copy**	**Edit, Create a Copy, Delete**
Message Analyzer Correlations	Yes	**New Union**, and **Manage Unions** items; **Import** and **Export** commands For more information about **Unions**, see Configuring and Managing Message Analyzer Unions.	No	**Create a Copy**	**Edit**, **Create a Copy**, **Delete**
Message Analyzer Chart View Layouts	Yes	**Save Current Layout As..**, **Save Current as Default User Layout**, **Load Default User Layout**, **Restore Application Default Layout**, **Manage Layouts**, **Edit**; **Import** and **Export** commands For more information on how to manage charts, see Managing Chart Viewer Layouts.	No	**Create a Copy**	**Edit**, **Create a Copy**, **Favorite**, **Delete**

ASSET COLLECTION LIBRARY NAME	LIBRARY DROP-DOWN LIST (Y/N)	MANAGEMENT ENTITIES	EXAMPLE ITEM (Y/N)	AVAILABLE CONTEXT	MENU COMMANDS
Message Analyzer Trace Scenarios	Yes	**Manage Trace Scenarios**; **Import** and **Export** commands For more information on how to manage trace scenarios, see Managing Trace Scenarios.	No	**Create a Copy**, **Favorite**	**Edit**, **Create a Copy**, **Favorite**, **Delete**
Message Analyzer Parsing Levels	Yes	None	No	None	None
Message Analyzer Window Layouts	Yes	None	No	None	None
Aliases	Yes	**Create Alias for <field> ...**, an **Analysis Grid** context menu item, the **Manage Aliases** item; **Import** and **Export** commands For more information about **Aliases**, see Using and Managing Message Analyzer Aliases. **Note:** The **Aliases** Library is not a collection that is updated by Microsoft, given that you typically customize aliases to your own environment.	No	None	**Edit**, **Create a Copy**, **Delete**

Exporting and Importing Library Items

In addition to the available commands specified in the previous table, the **Manage <*AssetType*>** dialog also enables you to perform the following operations:

- **Export Library items** — the **Manage <*AssetType*>** dialog enables you to create an export configuration by choosing the items you want to include in the export. For example, you can select individual items or all items in subcategories such as **Network** and **Examples** or you can select all items

in either or both top-level categories such as **Message Analyzer** or **My Items**, to comprise the export configuration. After you decide which items in an asset collection Library you want to export, click the **Export** button in the **Manage <*AssetType*>** dialog to display the **Save Library** dialog, which enables you to specify the following information:

- **Title** — specify a title for the Library items you are exporting.
- **Description** — optionally specify a description for the Library items you are exporting.
- **Author**— optionally specify the author of the Library items you are exporting.
- **Organization** — optionally specify the name of your organization.

When you finish specifying the **Save Library** dialog information, click **Save** to display the **Select Library Location** dialog from where you can save your export items as a *.asset file, with a file name of your choice, to a designated location such as a file share. Other team members can then access your saved collection and import it to the appropriate Library for use in their own Message Analyzer installations.

- **Import Library items** — the **Manage <*AssetType*>** dialog enables you to specify the items you want to import from a designated location by selecting the ones you want from the **Select Library to Open...** dialog. This dialog displays when you click the **Import** button in the **Manage <*AssetType*>** dialog. The **Select Library to Open...** dialog enables you to navigate to the share or other location that contains a collection of items that you want to import to an appropriate asset collection Library in your Message Analyzer installation. After you select the *.asset file you want to import and click **OK**, the **Select Items to Import** dialog displays from where you can specify the following:
  - **Select the items you want to import** — Message Analyzer automatically locates the appropriate Library in which to import selected items, based on the type of *.asset file you specified in the **Select Library to Open...** dialog, for example, a **My_Filters.asset** file for the **Message Analyzer Filters** asset collection Library. You can choose which items to import by selecting the check boxes for the items and/or categories you want.
  - **Choose the Library category for the imported items** — you can choose to place imported items into the existing categories by accepting the default **Use Existing Categories** option in the **Select Items to Import** dialog, or you can optionally deselect this check box and select the category from the **Add to Category** drop-down list. After you are finished with specifying the import configuration and you click **OK**, the **Manage <*AssetType*>** dialog appears with the selected items imported into the Library category that you specified.

> **NOTE**
> You are advised to not use the **Import** feature to retrieve asset collection items from a user-configured subscriber feed that is part of the Message Analyzer Sharing Infrastructure, as duplicates are not overwritten, but rather to obtain them by performing a download or auto-sync operation from the Message Analyzer **Asset Manager** dialog, as described in Managing Asset Collection Downloads and Updates.

Expanding User Libraries

There are four ways to add items to an asset collection Library, as follows:

- Create new items such as **Session Filters**, **Color Rules**, **Pattern Expressions**, **Aliases**, **Chart** viewer **Layouts**, and **Trace Scenarios**, or configure and save a new **Layout** for the **Analysis Grid** or **Grouping** viewer. After you create a new item, you can specify the Library category in which to place it.
- Modify an existing item and save it in the current user Library for that item.
- Import one or more items or asset collections from a user file share or other designated location.

- Download asset collections and updates from Microsoft when they are available.

> **NOTE**
>
> When you are using the common **Manage <*AssetType*>** dialog to **Export** asset collection Library items, several commands are available that you can utilize before you **Export**. Be aware that any changes that you make to an asset collection item are saved in the Library as soon as you make them. You can access these commands by right-clicking a Library item to select one of them, which can include those in the list below. Note that all of the commands in the list that follows are available only when you are selecting an item in the **My Items** category of an asset collection Library. For most other Libraries, only the **Create a Copy** and/or **Favorite** commands are available outside this category:

- **Edit**
- **Create a Copy**
- **Favorite**
- **Delete**

See the **Available Context Menu Commands** column in the previous table to review the commands that are available in the Library categories for each asset collection.

More Information

To learn more about how to manage downloads and updates for asset collection Libraries, see Managing Asset Collection Downloads and Updates.

See Also

User Libraries

Managing Asset Collection Downloads and Updates

2 minutes to read

The features that are available to manage asset collection downloads and updates consist of the following:

- Opting in or out of the auto-sync process for asset collections (and **OPN Parser** packages) at first Message Analyzer start up.
- Manually toggling the **Online**/**Offline** mode to moderate the update process.
- Syncing all displayed asset collections at once.
- Manually downloading specific asset collections.
- Manually setting asset collections for auto-synced updates.
- Searching and filtering asset collections.

These capabilities are described in the following topics:

Syncing Items on First Startup
Filtering and Searching For Items
Downloading Assets and Auto-Syncing Updates

More Information
To learn more about managing OPN Parser packages, see Managing Microsoft OPN Parser Packages.
To learn more about managing the Microsoft subscriber feed, see Managing the Default Subscriber Feed.

See Also

Asset Manager

Syncing Items on First Startup

2 minutes to read

This section describes the syncing options that you can utilize upon first Message Analyzer startup following installation. These options enable you to opt-in or opt-out of auto-syncing updates to your asset collection Libraries from a Microsoft web service. When you auto-sync a collection for updates, it is automatically downloaded to your Message Analyzer installation as new versions of existing collection items become available. This includes updates for asset collection Libraries, OPN Parser packages, and news items on the Message Analyzer **Start Page**.

Choosing to Opt In or Out of Auto-Synced Updates

When you first start Message Analyzer following your initial installation, you are presented with a **Welcome to Message Analyzer** dialog that prompts you to opt-in or out-out of auto-synced updates. The prompt indicates that while your PC is online, you can automatically receive updates to the previously indicated entities. The dialog options that are presented and the effect each one has consist of the following:

- **Update items** — if you select this option, the default asset collection Libraries are all set to the auto-sync state and moved to the **Settings** tab in the **Asset Manager** dialog, where the auto-sync state is indicated by a grey button containing circular arrows to the right of each asset collection. However, you can only receive updates that are pushed out by the web service if you have Message Analyzer set to the **Online** mode, which is the default setting in the **Asset Manager**. Moreover, if you toggle the **Online** button to the **Offline** mode, asset collections still retain the auto-sync status, but your local Libraries will not be updated until you reset to the **Online** mode.

- **Do not update items** — if you select this option, the default asset collection Libraries are not set for auto-syncing updates and they appear as un-synced on the **Downloads** tab of the **Asset Manager**, where the un-synced state is indicated by a server download status icon. However, even after choosing the **Do not update items** option, you can still set the default asset collection Libraries to auto-sync, either by clicking the **Sync All Displayed Items** button on the **Downloads** tab of the **Asset Manager** to auto-sync all collections, or by clicking the server download status icon of one or more individual collections, to display the **Item Download Options** dialog. From this dialog, you can select the **Automatically sync item updates when available** option to set the auto-sync state for any asset collection Library. Thereafter, any asset collection that you set for auto-syncing is moved to the **Settings** tab of the **Asset Manager**, where it displays the auto-synced status button.

See Also

Starting Message Analyzer for the First Time

Filtering and Searching For Items

2 minutes to read

The **Downloads** tab in the **Asset Manager** dialog displays all un-synced asset collection Libraries in a list, where the un-synced state is indicated by a server download status icon to the right of each collection. If the list of asset collections and **OPN Parser** packages becomes long and you want to quickly locate a particular collection, you can enter text in the search box that matches one or more characters or phrases in the name of the asset collection or parser package for which you are looking. If a match is found for the search text, the corresponding asset collection/s and parser package/s are displayed in the list of unsynced asset collections, while all other collections and parser packages are filtered out.

You can also filter for specific types of asset collections by selecting them by name from the **All asset types** drop-down list on the **Downloads** tab of the **Asset Manager**. For example, if you want to view only Filter collections on the **Downloads** tab, select **Filter** from the drop-down list. If you want to view only **OPN Parser** packages, select **OPN Package** from the drop-down list. To return to displaying all asset collections, which is the default setting, select **All asset types** from the drop-down list.

> **NOTE**
>
> When you auto-sync an asset collection Library, it is removed from the **Downloads** tab and appears on the **Settings** tab in the **Asset Manager** with the auto-sync status icon displaying to the right of that particular collection.

More Information

To learn more about auto-syncing asset collections, see Downloading Assets and Auto-Syncing Updates.

Downloading Assets and Auto-Syncing Updates

7 minutes to read

This section describes how to download the default asset collections and **OPN Parser** packages and how to configure them to automatically receive updates that are pushed out by a Microsoft web service. The manner in which you proceed in these tasks depends upon the update option you choose in the **Welcome to Microsoft Message Analyzer** dialog when you start Message Analyzer for the first time. If you opt-out of receiving updates by selecting the **Do not update items** option in the dialog, the latest version of the default asset collections for Message Analyzer are downloaded to your local user Libraries and none of the asset collections or parser packages are auto-synced to receive web service updates.

However, if you opt-in to receive updates by selecting the **Update items** option in the dialog, the latest version of the default asset collections are downloaded to your local user Libraries and all asset collections (and **OPN Parser** packages) are set to the auto-sync status so that your Message Analyzer installation can automatically receive the latest asset collection and parser package versions through the Microsoft web service, as they become available. The content that follows describes how to proceed with downloading asset collections and subscribing to automatic asset collection updates when you select either of the indicated startup options.

Opting Out of Automatic Updates

If you opted out of auto-synced updates for the default asset collections at first Message Analyzer startup, all asset collections and parser packages are initially downloaded and display on the **Downloads** tab of the **Asset Manager** dialog, with the server download status icon displayed to the right of each collection or package to indicate its un-synced state. In addition, the **Online**/**Offline** button is automatically set to the **Offline** mode to disable all collection updates. However, you still have the option to auto-sync any asset collection or parser package and perform downloads any time after first Message Analyzer startup. For example, even if you have not set an asset collection or parser package for auto-syncing updates, you can still perform a manual download of that collection or package to ensure you have the latest version; however, you cannot manually download again until you set the asset collection or parser package to the auto-sync state.

Auto-Syncing Asset Collections and OPN Parser Packages After Opt-out
After opting out of automatic updates, you can acquire the auto-sync state for asset collections and parser packages at any time. If you want to auto-sync these entities, do either of the following:

- **Auto-sync all asset collections and parser packages simultaneously** — click the **Sync All Displayed Items** button on the **Downloads** tab of the **Asset Manager**. This action configures all of the default asset collections and parser packages to automatically receive updates through the Microsoft web service as they become available. Also, all collections and packages are moved to the **Settings** tab of the **Asset Manager**, where the auto-synced status icon displays to the right of each collection or package. When an auto-sync update occurs, the corresponding update items in the **Message Analyzer** category of the appropriate local user asset collection Library or items in some **OPN Parser** package are refreshed to the latest version.

- **Auto-sync specific asset collections or OPN parser packages** — to manually set collections or packages for auto-syncing, click the server download status icon on the **Downloads** tab for a chosen collection or package to display the **Item Download Options** dialog, from where you can select the **Automatically sync item updates when available** option. After you select this option and click **OK** to exit the dialog, the following occurs:

    - The chosen asset collection or **OPN Parser** package will receive periodic updates, providing that you also have the **Online** mode set in the **Asset Manager**.

- The auto-synced asset collection or parser package is removed from the **Downloads** tab and reappears on the **Settings** tab of the **Asset Manager**, with a corresponding auto-sync status icon displaying to the right of the collection or package name and description.
- The **Message Analyzer** category of the local Library that corresponds to the asset collection is updated with the latest collection version and the cache that contains the parsers used by Message Analyzer is also updated. Thereafter, these items will continue to be periodically updated while they are set for auto-syncing.

Downloading Asset Collections and Parser Packages After Opt-out
If you opted out of automatic updates at first Message Analyzer startup, you can periodically download any asset collection or parser package by selecting the **Download once and don't automatically update** option in the **Item Download Options** dialog that displays when you click the status icon of any asset collection or parser package that is in the un-synced state. The following then occurs for the indicated entities:

- **Asset collections** — the **Select Items to Import** dialog displays, from where you can choose the items to import and the category in which to place them in the corresponding local asset collection Library. If you accept the default configuration, the asset collection items are placed in the **Examples** subcategory of the appropriate local Library. Thereafter, the chosen asset collection appears on the **Downloads** tab and the server download status icon continues to display to the right of that collection to indicate the un-synced state.
- **OPN Parser packages** — the **OPN package download** dialog displays and indicates that the package and its dependencies will no longer be downloaded and automatically installed. If you click **OK** to proceed, the parser package is removed from the **Settings** tab and appears on the **Downloads** tab in the un-synced state.

> **NOTE**
> When you employ the *download once* option for an asset collection or **OPN Parser** package, you are simply downloading the collection and opting out of the auto-sync process. However, you have the option to set the asset collection for auto-sync anytime thereafter, or to resume auto-sync after download of a previously synced item. After you choose the download option, the current version of the asset collection or parser package is downloaded from the Microsoft web service, even if the **Offline** mode is set.

Opting In to Automatic Updates

If you opted in to auto-sync updates at first Message Analyzer startup, then all of the default asset collections and parser packages are moved to the **Synced Collections** list and the **OPN Parsers** list, respectively, on the **Settings** tab of the **Asset Manager**, with the auto-sync status icon displayed to the right of each asset collection or parser package to indicate the synced state. In addition, the **Message Analyzer** category of each asset collection Library is updated with the latest collection version and the OPN parser cache is updated with the latest OPN parsers. Thereafter, updates to the asset collections and parser packages will be periodically pushed out by the Microsoft web service as they become available.

At that time, the asset collection items in the **Message Analyzer** category of the appropriate local user Libraries are refreshed with the latest versions and the OPN cache is updated as well. At any time, you can choose to download specific asset collections that are currently in the auto-synced state by clicking the auto-sync status icon for the collection on the **Settings** tab. Thereafter, the **Select Items to Import** dialog displays, from where you can choose the collection items to import and the Library category in which to place them, as previously indicated.

Working with Auto-Synced Asset Collections
If an asset collection is already synced and you click the auto-sync status icon for that collection on the **Settings**

tab of the **Asset Manager**, the **Import Autosynced Items** dialog displays and explains that the asset collection will no longer be synced and prompts you to import the current version to your local asset collection Library. If you click the **No** button to exit this dialog, the asset collection is removed from the **Settings** tab and reappears on the **Downloads** tab in the un-synced state, as indicated by the server download status icon for that collection. The asset collection items and the **Message Analyzer** category that contains them are also removed from the corresponding local Library for that asset collection. However, if you click the **Yes** button to exit the **Import Autosynced Items** dialog, the **Select Items to Import** dialog displays and enables you to choose the items to import and the category in which to place them in the corresponding local Library for the collection, as described earlier.

Working with Auto-Synced OPN Parser Packages

In the case of a parser package that is already synced, if you click the auto-sync status icon for the package on the **Settings** tab, the **OPN package update** dialog displays and indicates that the package and its dependencies will no longer be downloaded and automatically installed. If you click **OK** to proceed, the package reverts to the un-synced state, as indicated by the download status icon for the package on both the **Settings** and **Downloads** tab. Thereafter, you still have the option to reinstate the auto-synced state.

See Also

User Libraries
Managing User Libraries
Managing Microsoft OPN Parser Packages

Managing Microsoft OPN Parser Packages

7 minutes to read

By default, every Message Analyzer installation is provided with a baseline set of **OPN Parsers** that enable the PEF Runtime to decode messages that are captured by various Message Analyzer providers. These **OPN Parsers** are automatically copied to default locations during Message Analyzer installation. Thereafter, they are accessed by the PEF Runtime when parsers are required for decoding captured messages.

Downloading and Updating OPN Parsers

Message Analyzer enables you to download **OPN Parser** updates from a Microsoft web service that drives the **Message Analyzer** feed on the **Downloads** tab of the Message Analyzer **Asset Manager**. The **OPN Parsers** are listed on this tab and the **Settings** tab, from where you can manage your downloads with interactive status icons. As indicated earlier, the default behavior is to copy **OPN Parsers** to specific directory locations during Message Analyzer installation, but the default behavior does not automatically synchronize them for updates. In this state, all **OPN Parser** packages display server download icons that you can click for download options, either from the **Downloads** or **Settings** tab of the **Asset Manager**. To automatically synchronize for updates, you must set the *auto-sync* download option, which changes the **OPN Parser** download status; or you can elect to perform a single download without syncing for future updates.

> **TIP**
> You might want to perform a download after you install or upgrade Message Analyzer to ensure that you have the latest parser versions.

You can manage **OPN Parser** downloads and updates by selecting different options in the **Item Download Options** dialog, which displays when you click the server download status icon for an **OPN Parser** package. The options that you can select in this dialog along with the different actions that they invoke are described as follows:

- **Automatically sync item updates when available** — by selecting this option for a particular **OPN Parser** package, the latest version of that package is downloaded from the **Message Analyzer** feed so that you can use the latest parsers contained in the package for decoding messages that you capture with Message Analyzer. In addition, the **OPN Parser** package is configured for auto-synced updates, so that you automatically receive the latest parser versions as they become available in the future. Auto-syncing ensures that a selected parser package is always current and that Message Analyzer can parse messages from protocols that have undergone revisions to date. After an **OPN Parser** package is auto-synced, it is removed from the **Downloads** tab because it is automatically synchronized for future updates, and its status icon displays on the **Settings** tab as a grey button that contains circular arrows, to indicate the auto-synced state.
- **Download once and don't automatically update** — by selecting this option for an **OPN Parser** package, you opt out of the update process for the package. When this occurs, the current version of the package is downloaded from the **Message Analyzer** feed, but the package is not configured to auto-sync with the latest versions and download updates automatically. The **OPN Parser** package remains on the **Downloads** and **Settings** tabs and continues to display the server download status icon. However, you still have the option to auto-sync the package by clicking the server download status icon for the package on the **Downloads** tab to display the **Item Download Options** dialog, from where you can select the auto-sync option, even if you already downloaded the package once.

Other capabilities that you can employ when managing **OPN Parser** packages include the following:

- **Auto-sync all items** — by clicking the **Sync All Displayed Items** button on the **Downloads** tab of the **Asset Manager**, all **OPN Parser** packages are removed from the **Downloads** tab and appear on the **Settings** tab under the **OPN Parsers** list in the auto-synced state.
- **Reinstall or uninstall any OPN Parser package** — by clicking the appropriate icons on the **Settings** tab, such as the **Reinstall this OPN package** or **Uninstall this OPN package** icon, you can reinstall any **OPN Parser** package or uninstall one, respectively. Even if you temporarily uninstall an **OPN Parser** package, the feed and update design ensures that you can never lose an **OPN Parser** package and that you can always have access to the latest parser updates.
- **Work offline** — when you toggle the **Online** button on the **Downloads** page to the **Offline** state, you essentially prevent **OPN Parser** updates from being pushed out to your Message Analyzer installation from the web service. Going into the **Offline** state has no effect on your current **OPN Parser** download status, such that any auto-sync configuration that you previously set will automatically resume after you return to the **Online** state.

OPN Parser Groupings and Dependencies

The **OPN Parsers** are grouped into packages which contain components that relate to specific areas of functionality, such as applications, devices, communications, support, and so on. For example, the **Core Networking** package contains parsers for public network protocols and the **Microsoft Remote Desktop** package contains parsers for remote desktop application communications and management. In addition, some **OPN Parser** packages have dependencies on other **OPN Parser** packages, which means, for example, that syncing, downloading, or reinstalling one package will include any dependent packages as well. The dependencies between **OPN Parser** packages are automatically managed by Message Analyzer, such that the following occurs:

- When you download an **OPN Parser** package that depends on another **OPN Parser** package, it is automatically downloaded with the dependency package included.
- When you uninstall an **OPN Parser** package upon which other **OPN Parser** packages depend, you are prompted to also uninstall the dependent **OPN Parser** packages.
- When an **OPN Parser** package is set for auto-syncing updates, any **OPN Parser** packages upon which the synced package depends will be included in the updates.

OPN Parser Package Descriptions

The following table identifies and describes the baseline **OPN Parser** packages that are provided with Message Analyzer:

Table 28. OPN Parser Packages and Dependencies

NAME	DESCRIPTION	DEPENDENCY
Azure Storage Parsers Version 1.0	Contains parsers for both Azure Storage Analytics and client-side logs from Azure Storage Client Libraries	Core Networking
Core Networking Version 1.4	Contains parsers for public network protocols, upon which many other parsers depend. This is the basic foundation for protocol parsing functionality. Removing this group will prevent parsing of other major protocols.	Infrastructure

NAME	DESCRIPTION	DEPENDENCY
Device and Log File Version 1.4	Contains parsers for device traffic and text log files. Enable parsing of textlogs such as Cluster, IIS, Lync, Netlogon, SambaSysLogs, ULS, VMM, and so on.	None
Exchange ActiveSync Parsers Version 1.2	Contains parsers for Microsoft Exchange ActiveSync.	Microsoft Common, Core Networking, Infrastructure
Exchange MAPI Parsers Version 1.3	Contains parsers for Microsoft Exchange MAPI services.	Microsoft Common, Core Networking, Infrastructure
Exchange Web Services Parsers Version 1.2	Contains parsers for Microsoft Exchange Web Services protocols	Core Networking, Infrastructure
FSSHTTP and WOPI Parsers Version 1.2.1	Contains parsers for the FSSHTTP and WOPI protocols.	Microsoft Common, Core Networking, Infrastructure
Infrastructure Version 1.4.1	Contains infrastructure components that provide support functions, enumerations, and types for all other parsers and modules. This group does not contain actual parsers, but rather provides the basic foundation for protocol parsing. Removing this group may cause parsing functionality to fail.	None
Microsoft Common Version 1.4.1	Contains parsers for common Microsoft and public protocols, upon which many other Microsoft protocols depend.	Core Networking, Infrastructure
Microsoft File Sharing Version 1.4.1	Contains parsers for Microsoft Windows file sharing and branch cache protocols.	Microsoft Common, Core Networking, Infrastructure
Microsoft Identity and Security Version 1.4	Contains parsers for Microsoft Windows identity, authentication, authorization, and security protocols.	Microsoft Common, Core Networking, Infrastructure
Microsoft Others Version 1.4	Contains parsers for other, less common Microsoft protocols, to support special requirements.	Microsoft Common, Core Networking, Infrastructure
Microsoft Remote Desktop Version 1.4	Contains parsers for Microsoft Windows Remote Desktop communication and management protocols.	Microsoft Common, Core Networking, Infrastructure
Microsoft SMB2 Scenario Validation Version 1.4	Contains parsers that provide protocol scenario validation for SMB2 Negotiate, SessionSetup, TreeConnect, Logoff, and TreeDisconnect. Protocol scenario validation will check if the network traces violate defined message or field value ranges in protocol technical documents and report errors/warnings.	Microsoft Common, Core Networking, Infrastructure, File Sharing

NAME	DESCRIPTION	DEPENDENCY
Office and SharePoint Parsers Version 1.2	Contains parsers for the Microsoft Office and SharePoint protocols.	Microsoft Common, Core Networking, Infrastructure
Public Version 1.4.1	Contains parsers for public protocols that are not included in the Core Networking or Microsoft Common packages.	Microsoft Common, Core Networking, Infrastructure
Skype for Business Parsers Version 1.3	Contains parsers for Microsoft Skype services.	Core Networking, Public, Infrastructure

See Also

Managing Asset Collection Downloads and Updates
Asset Manager

Managing the Default Subscriber Feed

4 minutes to read

Message Analyzer has a default **Message Analyzer** subscriber feed that connects to a Microsoft web service, which provides asset collections and **OPN Parser** packages to your Message Analyzer installation. These collections and packages will be periodically updated over time as useful data manipulation, display, or tracing functionality and additional message parsers are developed at Microsoft for the community of Message Analyzer users. You have the option to auto-sync one or more asset collections or parser packages for updates and you can choose to perform a download of a particular collection or package. You can view these asset collections and parser packages in the **Asset Manager** dialog, synchronize with Microsoft updates, and download any of these items as needed.

Download and Update Processes

Downloads occur automatically for any asset collection or **OPN Parser** package that is set to the auto-sync state. An automatic asset collection download will update the **Message Analyzer** category items in the corresponding local user Library. An automatic **OPN Parser** download will update the local compilation cache that contains all message parsers.

Automatic downloads from the default **Message Analyzer** subscriber feed can occur only when the **Online**/**Offline** button is toggled to the **Online** mode. However, you can perform manual downloads even if the **Offline** mode is set. If you opted out of automatic updates when you started Message Analyzer for the first time, the asset collections and **OPN Parser** packages that are available for download are listed on the **Downloads** tab under the default **Message Analyzer** subscriber feed. From here, you can perform a download of current versions while the selected asset collections and/or **OPN Parser** packages remain in the un-synced state, or you can set any asset collection or **OPN Parser** package to auto-sync to receive automatic updates as they become available. Any collection or package that you set for auto-syncing updates is removed from the **Downloads** tab and appears on the **Settings** tab. This is also the case if you opted in for asset collection updates when you started Message Analyzer for the first time.

Feed Management Options

From the **Settings** tab, you can manage the default **Message Analyzer** subscriber feed, in addition to any custom feeds that you create. The options that are available for managing the default **Message Analyzer** feed consist of deleting and restoring it.

Deleting the Default Message Analyzer Feed

When you delete this feed, you are un-subscribing from it. During normal Message Analyzer operations, there is no reason why you should ever have to delete the **Message Analyzer** feed. However, if you do attempt to delete it by clicking the **X** to the right of the feed name, you are prompted that all auto-synced asset collections will be removed from your local Libraries, but that you can avoid such removal by clicking the auto-sync status icon of auto-synced collections and then importing the asset collection items into your local user Libraries. If you proceed with the deletion, all items in the **Message Analyzer** category of corresponding asset collection Libraries are removed. However, you can restore all items in this category for all your asset collection Libraries, restore all **OPN Parser** packages, and repopulate the collection and package lists on the **Asset Manager** dialog, by configuring a new feed from the **Settings** tab that points to the Microsoft web service.

Restoring the Default Message Analyzer Feed

To restore the **Message Analyzer** subscriber feed, click the **Add New Feed** button to display the **Add Feed Location** dialog. In the **Feed Name** text box, enter the text "Message Analyzer". In the **Location** text box, enter the URL of the default **Message Analyzer** feed as "https://go.microsoft.com/fwlink/?LinkID=401500". After you click **Add** to exit the dialog, all of the default asset collections and **OPN Parser** packages display on the **Downloads** tab under the default **Message Analyzer** feed with the server download status icon displaying to the right of each asset collection and parser package.

To restore all collections at once to the **Message Analyzer** category of your local user Libraries, click the **Sync All Displayed Items** on the **Downloads** tab.

> **NOTE**
> This action will also set all displayed **OPN Parser** packages to the auto-sync state. If there are any updates to **OPN Parser** packages after the default **Message Analyzer** feed is restored, all parser packages that are set to auto-sync will be automatically downloaded to your Message Analyzer installation when they are available.

To restore specific asset collections to the **Message Analyzer** category of your local user Libraries, click the server download status icon of selected asset collections and then select the **Automatically sync item updates when available** option in the **Item Download Options** dialog. If you select the **Download once and don't automatically update** option and click **OK** to exit the dialog, the **Select Items to Import** dialog displays, from where you can select the items you want to import to the associated Library and the category in which to place them. In this case, the asset collections will not be set for auto-sync and therefore will not be periodically and automatically updated.

Other Capabilities

Other management features for the default **Message Analyzer** subscriber feed enable you to do the following:

- Disable feed connections by toggling the **Online/Offline** button to the **Offline** mode. This action will block any update processes occurring in the background. This could be useful if you want to freeze your asset collection and/or **OPN Parser** package version configurations for consistency during protracted tracing and data analysis operations.
- Use the search box and/or filtering drop-down to locate specific items in the feed list on the **Downloads** tab, as described in Filtering and Searching For Items.

See Also

Asset Manager

Creating Custom User Feeds

3 minutes to read

Although Message Analyzer enables you to stay current with the latest asset collections and **OPN Parser** packages through the download and auto-syncing features, it also enables you to share asset collections directly with other users for collaboration and mutual use, as part of the Sharing Infrastructure capabilities. You can do this by configuring your own custom feeds. After you configure a user feed, you can export one or more items from any user Library as an asset collection and post it to a file share or other location to which the feed points. To perform an export, you must use the **Manage <*AssetType*>** dialog for the particular asset collection Library with which you are working. After you export an asset collection to a user feed, other team members can then acquire the collection by importing it into their local user Library for the specific collection type, by using the same **Manage <*AssetType*>** dialog for the particular Library. When exporting asset collection items, you can select specific collection items that you want to distribute to others, including any items that you have created or modified. When importing asset collection items, users can select which items they want to retrieve and the category in which to place them in their local user Library.

Configuring a User Feed

When you configure your own feed you must specify a feed **Location**, such as an SMB file share, web service, or other designated location, and you must also specify a **Feed Name**. The name you provide can be at your own discretion, but you might consider naming it in accordance with the team that will access the asset collection. To configure a user feed, go to the **Settings** tab in the **Asset Manager** dialog and click the **Add New Feed** button to display the **Add Feed Location** dialog. In the dialog, enter the name of the feed in the **Feed Name** text box and specify the file share path or other designated location in the **Location** text box. After you have entered this information, click the **Add** button to exit the dialog and create the new feed.

Note that you can create a feed and then place asset collections at the feed location; or you can place asset collections in a directory location that you intend to designate as the feed location, and then point to that location when you actually configure the **Location** value in the **Add Feed Location** dialog. However, if you use a file share as the location, keep in mind that you will need to configure the share with user permissions.

Working with a User Feed

After you configure a user feed, it displays on the **Downloads** tab of the **Asset Manager** in the feeds row, and on the **Settings** tab beneath the **Subscribed Feeds** label. When you click the user feed name on the **Downloads** tab, all the asset collections that exist on the feed are displayed in the **Downloads** list along with status icons to the right of each collection, just as it does for collections that are sourced from the default **Message Analyzer** subscriber feed. Thereafter, you can update existing collections or add others and make them available to team members or other users who subscribe to your feed. Other users can add your feed to their Message Analyzer installation by using the **Add New Feed** feature on their **Settings** tab. For others to subscribe to your feed, you will need to provide the feed location to them and ensure that they have permissions to the feed location, as appropriate. They can then view and download your asset collections by clicking the server download status icon for an asset collection and then selecting the download option to retrieve the collection for use in a particular local user Library.

However, for users to synchronize with asset collection *updates*, some manual configuration is necessary in the current Message Analyzer release, as described in Manual Item Update Synchronization. In future Message Analyzer releases, the Sharing Infrastructure publishing features will automatically enable others to synchronize to updates that you make to your asset collections on any user feed that you create.

See Also

Manual Item Update Synchronization

Manual Item Update Synchronization

2 minutes to read

This topic describes how to manually configure update synchronization for an asset collection that is published to a user-configured feed. In the Message Analyzer Sharing Infrastructure, the synchronization function of the publishing process relies upon consistency between the GUIDs of a published asset collection and the corresponding asset collection update version. This process occurs automatically whenever a default asset collection is set to the auto-sync state in the **Asset Manager** dialog. However, to successfully synchronize updates of asset collection items on a user-configured feed, some manual configuration is necessary to ensure that the GUID of the update asset collection is identical to the existing user asset collection on the user feed. To ensure that this occurs, it is incumbent upon the author of the collection update to manually assign the GUID of the existing user asset collection on the feed to the collection update version, so that users do not lose any items in their existing collection/s.

These GUIDs are specified in the *.metadata and *.asset file pair that exists for every user Library asset collection. You can manually configure these files for a collection update by copying the existing asset collection GUID into specific XML tags within these files, as described in the procedure that is referenced in Manually Configure Asset Collection Update Synchronization on a User Feed.

Sharing Asset Collections on a User File Share

2 minutes to read

You and other team members can directly share the items that are contained in your local asset collection Libraries. Message Analyzer provides a simple way to expose these assets so that you and others can retrieve them for collaboration and mutual use. To share asset collection items directly with others, you can use the **Export** command in the **Manage <*AssetType*>** dialog for the Library to save one or more items to a designated SMB file share or other location. In addition, you can use the **Import** command in the same dialog to access asset collection items that have been shared by others at the designated location. When exporting assets, you have the option to select specific collection items that you want to distribute to others, including any items that you have created or modified. When importing collection items, you can select which items you want to retrieve and the category in which to place them in an associated local asset collection Library.

> **NOTE**
> It is advisable to only import asset collections to your local Libraries through the Message Analyzer downloads and auto-sync features, because any duplicate items in a direct import will not overwrite existing items in your Libraries, which can result in duplicate Library items.

See Also

Creating Custom User Feeds
Managing User Libraries

Procedures: Using the Asset Management Features

9 minutes to read

This section contains procedures that demonstrate how to use the Message Analyzer asset collection download, auto-syncing, and sharing features for tasks that you are likely to perform on a consistent basis. Also included is a procedure that shows how to manually configure an asset collection for update synchronization.

Procedure Overviews
A brief description of each procedure is included here for review, as follows.

Dismiss Error Message - see below.

Download and Auto-Sync Asset Collections — provides an example of how to download an asset collection once without synchronizing it for updates; how to set a selected asset collection for auto-syncing updates; and how to synchronize all asset collections for updates.

Share Local Library Items on a File Share — provides an example of how to export and import user asset collection Library items for mutual sharing and collaboration with others.

Manually Configure Asset Collection Update Synchronization on a User Feed — shows how to manually configure an asset collection for update synchronization on a user-configured feed.

Dismiss Error message due to removal of the Microsoft Message Analyzer Feed service

This example demonstrates how to dismiss the error message that will appear when MMA is launched and attempts to connect to the back-end Feed service to check News and Assets updates after November 25 2019.

To dismiss this error message

1. Launch Microsoft Message Analyzer, in the **Tools** menu, select **Asset Manager**.
2. In the **Asset Manager** window, click on the 'Online' button to switch to 'Offline' state.
3. Close and restart Microsoft Message Analyzer. You will notice the error message disappears.

Download and Auto-Sync Asset Collections

This example demonstrates how to download default Library asset collections and how to auto-sync these collections for updates.

To download an asset collection

1. From the **Start** menu, **Start** page, or taskbar of your computer, click the **Microsoft Message Analyzer** icon to start Message Analyzer.
2. Open the **Asset Manager** dialog by clicking the global Message Analyzer **Tools** menu and then click the **Asset Manager** menu item.
3. In the **Asset Manager** dialog, click the **Downloads** tab (if it is not already selected) to display any asset collections that are not auto-synced, as indicated by the server download status icon to the right of the asset collection you want to download.
4. Click the server download status icon for an unsynced asset collection to display the **Item Download Options** dialog.

5. In the **Item Download Options** dialog, select the **Download once and don't automatically update** option and then click **OK** to exit the dialog and display the **Select Items to Import** dialog.
6. In the **Select Items to Import** dialog, accept the default selection of the **Use Existing Categories** check box, or optionally deselect this check box and specify the Library category in which to place asset collection items by clicking the **Add to Category** drop-down list and selecting a category.
7. In the **Select Items to Import** dialog, select the items and/or categories that contain the items you want to download.
8. Click **OK** to exit the **Select Items to Import** dialog.
9. Open the local Library that corresponds to the type of asset collection you downloaded and observe that the Library is populated with the items you selected and in the category you specified.

To auto-sync an asset collection

1. From the **Start** menu, **Start** page, or taskbar of your computer, click the **Microsoft Message Analyzer** icon to start Message Analyzer.
2. Click the **Downloads** tab in the **Asset Manager** dialog to display any asset collections that are not auto-synced, as indicated by the server download status icon to the right of a target asset collection.
3. Click the server download status icon for the asset collection that you want to auto-sync for updates, to display the **Item Download Options** dialog.
4. In the **Item Download Options** dialog, select the **Automatically sync item updates when available** option and then click **OK** to exit the dialog.

   The asset collection that you auto-synced is removed from the **Downloads** tab and reappears on the **Settings** tab with the auto-sync status icon displaying to the right of the collection. The auto-synced state indicates that the asset collection in your Message Analyzer installation will be periodically and automatically refreshed with updates as they become available.
5. To auto-sync all asset collections that are currently displayed on the **Downloads** tab in the **Asset Manager**, click the **Sync All Displayed Items** button.

   The asset collections that you auto-synced are removed from the **Downloads** tab and reappear on the **Settings** tab with the auto-sync status icon displaying to the right of each collection.

> **NOTE**
> When you click the **Sync All Displayed Items** button on the **Downloads** tab to set all asset collections to the auto-sync state, this includes all **OPN Parser** packages.

Share Local Library Items on a File Share

This example shows how to export **Filter** items as an asset collection from the centralized Filter Expression **Library** to a user-configured file share or other designated location to share these items directly with other users. This example also specifies how users can retrieve the asset collections that you post to such a designated location.

To export a Filter asset collection

1. From the **Start** menu, **Start** page, or taskbar of your computer, click the **Microsoft Message Analyzer** icon to start Message Analyzer.
2. Optionally, load a saved trace file into Message Analyzer through a Data Retrieval Session.
3. On the default Filter panel of the Filtering toolbar above the main analysis surface of Message Analyzer, click the **Library** drop-down list to expose the asset collection Library items and management features.

4. In the **Library** drop-down list, click the **Manage Filters** item to open the **Manage Filter** dialog.
5. In the **Manage Filter** dialog, select the collection items and/or **Library** categories that contain the items you want to export for sharing by placing a check mark in the appropriate check boxes.
6. Click the **Export** button on the toolbar of the **Manage Filter** dialog to open the **Save Library** dialog.
7. In the **Save Library** dialog, enter a name for the asset collection in the **Title** text box and the collection author in the **Author** text box. Optionally, add **Description** and **Organization** information.
8. Click the **Save** button to exit the **Save Library** dialog and to display the **Select Library Location...** dialog, from where you can navigate to the file share location where you intend to post the asset collection.
9. After you specify a **File name** for the asset collection, click the **Save** button to exit the **Select Library Location...** dialog, at which time your asset collection is posted to the file share location.

> **NOTE**
> Ensure that users have your file share location information and appropriate permissions to access the share or other location.

To import a Filter asset collection

1. From the **Start** menu, **Start** page, or taskbar of your computer, click the **Microsoft Message Analyzer** icon to start Message Analyzer.
2. Optionally, load a saved trace file into Message Analyzer through a Data Retrieval Session.
3. On the default Filter panel of the Filtering toolbar above the main analysis surface of Message Analyzer, click the **Library** drop-down list to expose the asset collection Library items and management features.
4. In the **Library** drop-down list, click the **Manage Filters** item to open the **Manage Filter** dialog.
5. In the **Manage Filter** dialog, click the **Import** button on the toolbar to display the **Select Library to Open...** dialog, from where you can navigate to the file share that contains the shared asset collection.
6. Select the *.asset file that corresponds to the target asset collection and then click the **Open** button to exit the **Select Library to Open** dialog and to open the **Select Items to Import** dialog.
7. In the **Select Items to Import** dialog, accept the default selection of the **Use Existing Categories** check box, or optionally deselect this check box and specify the Library category in which to place collection items, by clicking the **Add to Category** drop-down list and selecting a category.
8. In the **Select Items to Import** dialog, select the collection items and/or categories that contain the items you want to import.
9. Click **OK** to exit the **Select Items to Import** dialog and populate your centralized Filter Expression **Library** with the items you selected.

Manually Configure Asset Collection Update Synchronization on a User Feed

This example describes the process that you must follow to manually synchronize an asset collection for updates on a user-configured feed that points to a file share or other location where the collection is posted.

To perform manual asset update synchronization

1. From the **Start** menu, **Start** page, or taskbar of your computer, click the **Microsoft Message Analyzer** icon to start Message Analyzer.

2. In the **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu, click the **Settings** tab to display the list of **Subscribed Feeds**.
3. Click the link under the feed name containing the asset collection that you want to update, to display the *.asset and *.metadata files for the collection that exists at the feed location.
4. For the asset collection that you intend to update, right-click the *.metadata file and select the **Open With** context menu item so that you can specify an XML editor or the Visual Studio application to open the file.
5. In the opened metadata file, copy the GUID in the <UniqueId> or <GroupId> tag and store it in a temporary location.
6. Close the file and exit the editor application.
7. In the Message Analyzer user interface, locate the Library that contains the asset collection items you want to share with other users on your feed, click the Library drop-down list, and then select the **Manage *<AssetType>*** item.
8. In the **Manage *<AssetType>*** dialog, select the asset collection items you want to include in the update and then click **Export** to display the **Save Library** dialog.
9. In the **Save Library** dialog, specify values for **Title** and **Author**, and optionally provide a **Description** and **Organization** information.
10. Click **Save** to exit the **Save Library** dialog and to display the **Select Library Location...** dialog, from where you can navigate to the feed location. Note that you will be overwriting an existing collection when you click **Save**.
11. Enter the asset collection **Title** that you specified in the **Save Library** dialog as the **File name** for the collection in the **Select Library Location...** dialog.
12. Click **Save** to exit the **Select Library Location...** dialog.
13. Click **Close** to exit the **Manage *<AssetType>*** dialog.
14. On the **Settings** tab in the **Asset Manager** dialog, click the link beneath the feed name that will contain your updated asset collection.
15. Right-click the *.metadata file and *.asset file for the update collection, and then do the following:
    - Select the **Open** command to open the files.
    - Paste the GUID that you obtained from step 5 into the <UniqueId/> and <GroupId/> tags in each file.
    - Incrementally increase the <Revision/> tag value in each file with a new matched value to indicate that a revision is available to users.
    - Save the changes and close each file.
16. Inform users that you have an update that they can download if they are subscribed to your feed, which should appear on the **Downloads** tab in the **Asset Manager** dialog of the user's Message Analyzer installation.

Users should then be able to use the download or auto-sync option in the **Item Download Options** dialog to update their local Library with the latest version of your asset collection. After users update a Library, the collection items should appear in the Library under a category name that matches the feed name.

See Also

Extending Message Analyzer Data Viewing Capabilities

3 minutes to read

Although Message Analyzer provides many built-in data viewers that attempt to address common analysis scenarios, it can be to your advantage to customize your data analysis environment with additional data viewing capabilities that meet your specific needs. Message Analyzer accommodates for these requirements by allowing you to extend Message Analyzer data viewing capabilities with custom **Layouts** of your own design that you can create for the **Chart** viewer. When you create your own **Layout**, you can configure any one of several types of graphic data visualizer components along with one or more data formulas that you can specifically craft to provide data analysis capabilities that streamline your work.

> **IMPORTANT**
>
> The process of creating **Chart** viewer **Layouts** has been simplified in Microsoft Message Analyzer v1.4 to make it easier for users to create them. One of the most apparent departures from the former process that simplifies the configuration is the fact that you can now add one visualizer component only to any one **Layout**, as indicated immediately below. The process to create formulas has been simplified as well, as described later in this section.

When you are creating a new **Chart** viewer **Layout**, you have the option to select one of the following types of graphic visualizer components for your **Layout**. Note that the built-in **Layouts** for the **Chart** viewer each use one of these same data visualizer components, as follows:

- **Bar** element
- **Pie** slice
- **Timeline** graph
- **Table** grid

What You Will Learn

In the topics of this section, you will learn about the functions of **Chart Layout** configuration features, how to use them to create new **Layouts** that you can display whenever you require them for analysis, and how to manage and share your custom **Layout** assets with others, through the Message Analyzer Sharing Infrastructure. A walkthrough of a built-in **Chart** viewer **Layout** that Message Analyzer provides by default is included for the benefit of exposing the behind-the-scenes configuration of a working **Layout** so that you can come up to speed on the process of creating a new **Chart** viewer **Layout**.

In This Section

Configuring Chart Viewer Layouts — learn about the constraints with which you must work when creating a **Chart** viewer **Layout** of your own design, the built-in **Chart** viewer **Layouts** that are provided with Message Analyzer by default, the visualizer components that Message Analyzer uses in the built-in **Chart** viewer **Layouts** and that you can likewise use in your own **Layouts**, along with the criteria for choosing a visualizer component type for a custom **Layout** of your own. You can also read an overview about configuring a custom **Chart** viewer **Layout**.

Using the Edit Chart Layout Dialog — learn about the controls and features that you can use to create and

save your own custom **Layouts** with a visualizer component that you specify. Also describes how to use the **Edit Chart Layout** dialog to set **Chart Properties**, **Series Fields**, and **Values**; and the **Formula Editor** dialog in which you can create data formulas based on a specified operation that manipulates message field data or computed values.

Configuration Walkthrough of a Built-In Chart Viewer Layout — perform a walkthrough of the built-in **TCP/UDP Conversations by Message Count Layout** for the **Chart** viewer to familiarize yourself with the configuration features that you can use to create a functioning **Layout**. This particular **Layout** exposes the network conversations in a set of trace results and the transports that carried those conversations, along with a set of statistics for analyzing performance that includes message volume, payload volume in bytes, data transmission rate, and duration for each conversation.

Managing Chart Viewer Layouts — learn how you can use the Message Analyzer Sharing Infrastructure to manage the **Message Analyzer Chart View Layouts** Library asset collection, which includes managing individual collection items through context menu commands such as **Edit**, **Create a Copy**, and **Delete**; and creating an item collection that you can share with others or retrieving a collection that was shared by others.

Configuring Chart Viewer Layouts

11 minutes to read

The built-in **Chart** viewer **Layouts** that are provided with Message Analyzer by default represent various ways to organize the display of data using any of the four data visualizer components that are available for **Layouts**. When configuring a custom **Layout** of your own design, the following constraints apply:

- You can use only one of the four types of graphic visualizer components in any one **Layout**.
- You can configure one or more data formulas for each **Layout**.
- You can only display data in your **Layout** based on message fields that are available from the **Field Chooser Tool Window**, in addition to global **Annotations** and **Properties**.
- Any new **Layout** that you create must be modified from an existing **Layout** and saved with new metadata that you define, such as the **Layout** name, description, and category.

This section begins with a list of the built-in **Chart** viewer **Layouts** that are included with every Message Analyzer installation. You can edit any one of these **Layouts** and save it as a new custom **Layout** of your own that will appear in the **My Items** category of the **Chart**/**Layout** drop-down list that is accessible from the global Message Analyzer **Session** menu. You might navigate to the topics in the lists and review some of the built-in **Layout** features to get an idea of the types of data displays that you can create and how you can use them in analysis. This section also provides a brief description of the graphic visualizer components that you can use in any custom **Layout** that you create, some background on how to choose a visualizer component for a custom **Chart** viewer **Layout** of your own, and a brief overview of configuring a custom **Layout**.

Built-In Chart Viewer Layouts

Message Analyzer provides numerous **Layouts** for the **Chart** viewer that are accessible from the locations that follow. Note that each **Layout** in the drop-down lists that are described immediately below has an icon to the left of the **Layout** name that identifies the type of graphic visualizer component that it uses to display data.

- The **New Viewer** drop-down list on the global Message Analyzer toolbar. You can display an uncategorized list of **Layouts** by clicking the **Chart** item in this drop-down list.
- The **Session Explorer** context menu, which is accessible by right-clicking anywhere in the **Session Explorer Tool Window**. You can then select the **New Viewer** item in the context menu, and in turn select **Charts** to display an uncategorized list of **Layouts**.
- In the **Layout** drop-down list that is accessible from the global Message Analyzer **Session** menu. You can display the fully categorized list of **Layouts** by clicking the **Chart** item in the global Message Analyzer **Session** menu and then selecting **Layout**.

> **NOTE**
>
> The **Start With** drop-down list in the **New Session** dialog enables you to select a default **Chart** item for session startup that utilizes a **Bar** element visualizer component. See Selecting a Session Data Viewer for further details.

Subcategories for the Built-In Chart Viewer Layouts

The **Layout** list that you can access from the **Session** menu, as previously described, is organized into a top-level **Message Analyzer** category with various subcategories in which the **Chart** viewer **Layouts** appear. The subcategories in the following list contain the **Layout** names, which are each linked to the corresponding topic so

that you review their functionality and analysis capabilities:

- **HTTP** subcategory **Layouts**:
  - HTTP Content Type Payloads
  - HTTP Content Type Volumes
- **General** subcategory **Layouts**
  - Average Elapsed Time for Operations
  - Average Response Time for Operations
  - Cluster Levels
  - Event Log IDs
  - IP/Ethernet Conversations by Message Count
  - IP Ethernet Conversations by Message Count Top 20
  - TCP/UDP Conversations by Message Count
  - TCP/UDP Conversations by Message Count Top 20
  - Top Level Protocols Message Count
  - Top Level Protocols Message Count Over Time
- **Network** subcategory **Layouts**
  - IIS Log HTTP Traffic Volumes
  - IIS Log Server Bytes by Host over Time
  - IIS Log Top URI Bytes
  - IIS Log Top URIs by Time
  - TCP Rate and Diagnosis
  - TCP Stevens Graph
  - Top Talkers
  - Top Talkers Top 20
- **NetLogon** subcategory **Layouts**
  - Netlogon Message Types
- **Networking** subcategory **Layouts**
  - NTP Time Offset
- **Common** subcategory **Layouts**
  - Perfmon Log (.blg)
- **File Sharing** subcategory **Layouts**:
  - SMB File Stats
  - SMB Reads and Writes Bytes Sent
  - SMB Reads and Writes Bytes/Second

- SMB/SMB2 Service Performance
- SMB Top Commands
- SMB Top Talkers
- SysLog Levels

Graphic Visualizer Components

Each of the built-in **Layouts** in the previous list use only one of the four different types of graphic data visualizer components that are available. You can also make use of any one of these components when you are configuring a new **Layout** that organizes and presents data in a unique format, which can include rendering top-level data summaries in **Bar** element, **Pie** slice, event **Timeline** graph, and **Table** format, as described in Data Viewer Concepts. Note that you can see all four of the available visualizer components in use if you display the Protocol Dashboard viewer from the **Charts (Deprecated)** drop-down list, which is accessible from the **New Viewer** drop-down list on the global Message Analyzer toolbar.

Many of the **Chart** viewer **Layouts** are intended to work together with other data viewers to create an integrated and interactive analysis environment, as described in Working With Message Analyzer Profiles. The interactive features of a **Layout** are initiated when you double-click the elements of a visualizer component that represents some message quantity or other value, for example, message volume, payload length, or some other field value. This action results in interactively driving the display of the element messages into another viewer such as the **Analysis Grid** to create an analysis context that focuses on the element messages only. Note that all of the built-in **Chart** viewer **Layouts** are capable of interactively driving the display of messages in other viewers, whether or not they are configured in a Message Analyzer **Profile** which can be enabled or disabled.

The list that follows provides a brief description of the four types of visualizer components and the elements they contain that can interact with the **Analysis Grid** viewer and other viewers by double-clicking them, or in some instances by using a single-click.

- **Bar** element — contains bar elements that each represent a group of captured messages that have a common field, property, message type, or other entity to which a data manipulation formula has been applied. Double-click any single bar element to view the associated messages in a new instance of the **Analysis Grid** viewer.
- **Pie** slice — divided into slices or section elements that each represent a group of captured messages that have a common field, property, message type, or other entity to which a data manipulation formula has been applied. Double-click any single slice to view the associated messages in a new instance of the **Analysis Grid** viewer.
- **Timeline** graph — displays an interconnected timeline element across trace boundaries for common properties, fields, values, or other entities that are contained in captured messages, to expose the points in time on the X-axis where those entities were sent and/or received, versus other values on the Y-axis that might represent the application of a particular formula, such as the **Count** of identical values for a particular field. Double-click a single timeline element or node to view the associated messages in a new instance of the **Analysis Grid** viewer.

> **NOTE**
> The **Timeline** visualizer is enabled for zooming into chosen windows of time.

- **Table** grid — contains data row and column elements that you can organize to correlate values with entities such as message fields and properties, or formulas based on field and properties, for example the **Average** of a set of field values. Double-click a data row to view the associated messages in a new instance of the **Analysis**

Grid viewer.

Choosing a Visualizer Component for a Custom Layout

When you are creating a custom **Chart** viewer **Layout** of your own design, you should consider several factors before you select the visualizer component that you want to work with, as follows:

- **Troubleshooting context** — this refers to the environment or circumstances in which you typically expect to experience issues, for example, connectivity, performance, security, diagnostics, Internet, and so on.

  For example, if you are having performance issues, a high-level view of summary data with a **Bar** element or **Pie** slice visualizer component might be the best choice to obtain a quick assessment of specific types of performance data. Alternatively, a **Table** grid visualizer component can also work when you want to display a combination of parsed field data and computed statistical values in tabular format.

- **Message type** — this is the type of message data that you are working with. Generally, the messages that are issued by the particular protocol or module with which you are working point to the types of information that you can expose for analysis.

  For example, this could be the message packets that are issued by HTTP, TCP, or LDAP, and even ETW layer events.

- **Information to expose** — consists of the aspects of message data that you want to expose to optimize the analysis perspective that you can obtain from the displayed results in your **Layout**.

  For example, if you are interested in working with IPv4 messages, you might want to expose network conversations with the **IPv4.Datagram.Network** field in your **Layout** along with their associated message volumes across a set of trace results. If you are interested in working with SMB2 messages, you might want to expose the average response time for all messages associated with the SMB2 queries that occurred in a set of trace results. The average response time would be calculated by a **Layout** formula that is based on the **ResponseTime** Global Annotation; see SMB/SMB2 Service Performance for further information.

> **TIP**
> You have the option to employ **Unions** in any **Layout** that you modify, which includes **Union** sets.

- **Data presentation format** — the format in which you present data should align with the level of detail that you want to see, for example, a high-level summary or low-level details that include message field data and other computed values, or events occurring in the context of time, as follows:

  - **High-level summary data** — enables you to see top-level information that you can assess at-a-glance. A **Bar** element or **Pie** slice visualizer component would be a good choice for this type of data presentation because they provide a graphic format that displays the relative distribution of specified values that you can set in ascending or descending order, where each **Bar** element length or **Pie** slice size represents a particular volume of such values, for example message count or cumulative byte volume.

  - **Low-level details data** — enables you to review significant message details and computed values that provide a set of statistics that can expose the cause of various issues and failures. A **Table** grid visualizer component would be a good choice for this type of data presentation to expose many different field values and calculated values that result in a statistical view of your data.

  - **Event data** — enables you to assess events that occur in time. A **Timeline** visualizer component would be a good choice for this type of data presentation, given that you can visually identify the points in time where any particular message occurred in a set of trace results. The **Timeline**

component also enables you to use presets or configured windows of time to drill down into any chosen time slot for detailed analysis.

Note that you can link to some of the built-in **Chart** viewer **Layout** descriptions in the above section "Subcategories for the Built-In Chart Viewer Layouts" to understand chosen **Layouts**, and then display them in an Analysis Session to see the various types of graphic visualizer components in action. To locate the **Chart** viewer **Layouts** in Message Analyzer, see the previous Built-In Chart Viewer Layouts topic.

Overview of Configuring a Custom Chart Viewer Layout

Message Analyzer enables you to create your own custom **Layouts** containing one of the four types of data visualizer components. The protocol/module types, fields, properties, and the formulas that you can apply to these entities can provide a quick overview of trace activity at-a-glance or a detailed analysis of statistical data, to enhance your data analysis perspectives. You can use formulas to manipulate the values of fields, properties, and other entities in the messages that you capture to create unique data representations. An example of this is the **Bar** element visualizer component that is used in the **SMB Top Commands Layout**. This data visualizer depicts the relative distribution of traffic volume, from the highest to the lowest volume, for SMB commands in a set of trace results. This in turn can help you to quickly evaluate the SMB commands that are consuming the most bandwidth, which may point to other issues.

Editing a Built-In Chart Viewer Layout

If you want to create a custom **Layout**, you will need to edit an existing **Layout** first and then save your changes under a specified **Layout** name. To do so, you must first display the **Layout** you want to edit by selecting it in any of the locations specified in the Built-In Chart Viewer Layouts section. Thereafter, to begin editing, you will need to select the **Chart** item in the global Message Analyzer **Session** menu and then select the **Edit** item in the **Chart** drop-down list to display the **Edit Chart Layout** dialog. You can then make use of the controls in the **Edit Chart Layout** and **Formula Editor** dialogs to modify the current **Layout**.

> **IMPORTANT**
>
> Editing an existing **Layout** and saving it under a different name is the only way you can now create your own **Layouts**, as the former **New Chart** command and others are no longer available. This change is advantageous given that it streamlines the UI and results in fewer clicks and selections to create a custom **Layout**.

In the topic that follows, a link to which is given immediately below, you will learn how to use the controls of the **Edit Chart Layout** dialog to edit an existing **Layout**. After you complete your modifications, you will then learn how to save your changes as a custom **Layout** of your own design.

Using the Edit Chart Layout Dialog

Using the Edit Chart Layout Dialog

15 minutes to read

To obtain access to the configuration controls and features that you will need to create a custom **Chart** viewer **Layout**, you must have a **Layout** already displaying in an Analysis Session, for example, Data Retrieval Session or Live Trace Session results. When you have a **Layout** displayed, the drop-down lists and commands required to edit and save a custom **Layout** of your own design are enabled; otherwise, such lists and commands are unavailable. The required drop-downs and commands for editing and saving a **Layout** appear in the global Message Analyzer **Session** menu, for example, the **Chart** item which displays a drop-down list that contains the **Edit** command and the **Layout** item from where you can save a modified **Layout**.

The sections that immediately follow describe how to use the configuration controls of the **Edit Chart Layout** to create your own custom **Layout** configuration. You can then learn how to save the **Layout** with a new **Name**, **Description**, and **Category** when complete, as described in Saving a Custom Chart Viewer Layout.

Editing an Existing Chart Viewer Layout

If you have a **Chart** viewer **Layout** already in focus and displaying data from the current set of trace results, you can edit the **Layout** and save your changes. To begin, click the **Edit** item on the toolbar that appears above the **Chart** analysis surface to display the **Edit Chart Layout** dialog, from where you can modify the currently displayed **Layout** as needed for your own custom design. Alternatively, you can click the **Edit** item in the **Chart** drop-down list on the global Message Analyzer **Session** menu to display this dialog, which is shown in the figure that follows.

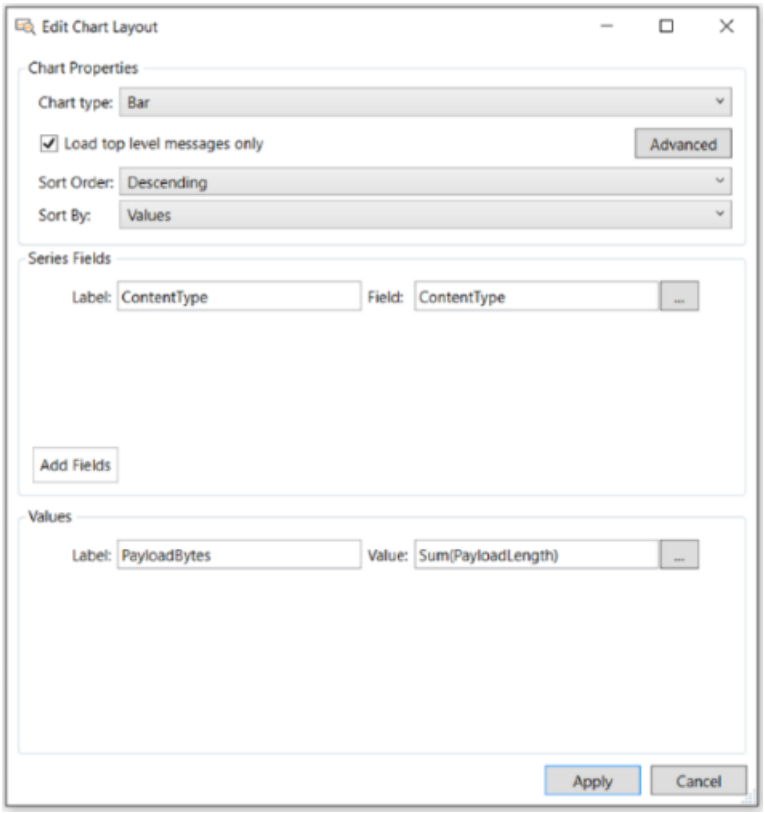

Figure 72: Message Analyzer Edit Chart Layout dialog

The controls of the **Edit Chart Layout** dialog that you will use to edit a **Chart** viewer **Layout** are described in the list that follows. The dialog contains the following three panes of controls and data entry selectors, as shown in the figure:

- **Chart Properties** pane
  - **Chart type** — a drop-down list that enables you to select the type of visualizer component you want to use in your **Layout**. You can select one of the four types that are available, which includes **Bar** element, **Pie** slice, **Timeline** graph, and **Table** grid components.
  - **Load top-level messages only** — a check box that enables you to specify whether the visualizer component will display top-level messages only. If this check box is *unselected*, the visualizer component will display data for all messages, including top-level and origins (stack) messages. Note that this setting can affect performance since more messages will have to be processed. If this check box is *selected*, the visualizer component will display data from the top-level messages only, which improves performance by processing fewer messages. For some context, the **Analysis Grid** viewer displays top-level message nodes by default.
  - **Advanced** button — this control is available for the **Bar**, **Pie**, and **Timeline** components only. The list immediately below describes the features of the **Generic Dialog Window** that displays when you click the **Advanced** button while the specified visualizer component is selected in the **Chart type** drop-down list:
    - **Bar** — clicking the **Advanced** button while this visualizer component is selected in the **Chart type** drop-down displays the **Generic Dialog Window**, from where you can use the following features to achieve the indicated results:
      - **Maximum number of items shown** text box — enables you to specify the maximum number of bar elements to display in your **Layout**, in which specific data values will be represented. You can manually type a number in the text box, or you can use the up-down control to the right of the text box to automate your input.
      - **Hide Others Category** check box — select this check box to enable your visualizer component to ignore the "Other" category that often displays when values for entities such as messages, fields, properties, and so on, far exceed the values of other visualizer component elements. When this occurs, **Other** category volumes can be so large as to overwhelm the display of other values of interest in your **Layout** component. However, you can unselect this check box if you want to enable the display of the **Other** category.
    - **Pie** — clicking the **Advanced** button while this visualizer component is selected in the **Chart type** drop-down displays the **Generic Dialog Window**, from where you can use the identical features that exist for the **Bar** element visualizer component, as previously described.
    - **Timeline** — clicking the **Advanced** button while this visualizer component is selected in the **Chart type** drop-down displays the **Generic Dialog Window**, from where you can use the following features to achieve the indicated results:
      - **Maximum number of items shown** text box — enables you to specify the maximum number of legend items that will be selected by default, which in turn enables the corresponding message lines in the **Timeline** visualizer of your **Layout** to display, in which specific data values will be represented. You can manually type a number in the text box, or you can use the up-down control to the right of the text box to automate your input.
      - **Data Aggregation Mode** drop-down selector — works with the bucketization feature to define how data will display in a bucketized data point or node. The following values are provided for selection:
        - **Add values** — adds all values in a node together; the resulting value is visible when your mouse hovers over a bucketized data point.

        - **Average values** — takes the average of all values in a node; the resulting value is visible when your mouse hovers over a bucketized data point.
        - **List values** — lists all values in a node; all values are visible when your mouse hovers over a bucketized data point.
    - **DisableBucketization** check box — by default, the display of screen values are bucketized according to a time resolution algorithm that sets a default number of values buckets that can appear in a given distance on the visualizer component screen real estate. Based on the size of the screen, a calculation is made to allow the visualizer component to contain one or more values in a bucket, for example, a node on the **Timeline** graph. If you hover over such a node with the mouse, Message Analyzer will display a pop up that provides messaging information for the particular hovered-over data point. You can disable this process by unselecting this check box.
    - **Value Number Format** — enables you to specify the number format for the Y-axis in a **Timeline** visualizer component. The default value is "N", for the Number format.

      For more information about possible settings for this property, see Standard Numeric Format Strings in the MSDN Library.
- **Sort Order** drop-down selector — enables you to specify the order in which a data field value (Series) or formula result (Value) is sorted in a visualizer component. You can select from the following three options.
    - **Ascending** — typically sorts from the lowest to highest field values.
    - **Descending** — typically sorts from the highest to lowest field values.
    - **Unspecified** — field data follows the order in which messages were captured.
- **Sort By** drop-down selector — enables you to configure what the specified **Sort Order** will apply to. You can select from the following two options:
    - **Series** — causes the specified **Sort Order** to apply to a configured data field value, such as **PayloadLength** in bytes.
    - **Values** — causes the specified **Sort Order** to apply to a configured formula output value, such as the result of an **Average**, **Count of Occurrences**, or **Cumulative Addition** operation upon a field, or other operations upon a set of fields.

- **Series Fields** pane — provides controls that enable you to add one or more sets of field controls that each consist of a **Label** text box and a read-only **Field** text box. One set of controls displays by default, although you can add more control sets by clicking the **Add Fields** button. For each set, you can specify a **Field** name for which you want data to display in your **Layout** and a **Label** that describes the data field contents. Each field that you define becomes part of the *series* of fields that you configure. For example, multiple fields in the **Series Fields** pane will be separated by commas in a **Bar** element visualizer, or will become column names in a **Table** grid visualizer.

> **NOTE**
> You will typically use the **Series Fields** pane controls to configure **Fields** that inherently display a value by default, for example, an IP conversation (**Network** field) or bytes value (**PayloadLength** field). On the other hand, you will typically use the **Values** pane controls to configure formulas that use a **Computed Value** argument along with built-in operations that you can select to manipulate one or more data field values and create statistics that are useful for the type of analysis you are performing. However, in the **Values** pane, you can also create formulas that use a **Message Field** argument that works with a specified operation upon one selected field in **Field Chooser Tool Window**.

The following describes the controls in the **Series Fields** pane of the **Edit Chart Layout** dialog:

- **Label** — a text box in which you to specify a name for the series of fields that will display in the **Layout**. Currently, this feature is most useful with the **Table** grid visualizer component, where the **Label** value that you specify for each field added to the **Series Field** pane displays as a table column name in your **Layout**.
- **Field** — a read-only text box that specifies the name of a selected field for which data will display in your **Layout**. You can select a field by clicking the ellipsis (**...**) button to the right of the read-only text box and then locating and double-clicking a field name in the **Field Chooser** window.
- **Ellipsis** — a button that enables you to display the **Field Chooser** by clicking anywhere on the button.
- **Add Fields** — a button that adds additional sets of field controls to the **Series Fields** pane when clicked.

- **Values** pane — provides controls that enable you to create formulas that manipulate values of chosen field/s, to provide statistics that enhance the functionality of a **Layout**. For example, if you selected HTTP **ContentType** as a field in the **Series Fields** pane, you might want to also see the sum of all **PayloadLength** values from all the HTTP responses that are associated with each **ContentType**, in order to assess the loads that a web server is delivering to an HTTP client.

  Therefore, in this case, you would specify the HTTP response **PayloadLength** field in the **Argument value** text box of the **Formula Editor** dialog and you would also select the **Cumulative addition** item in **Formula** drop-down list of this dialog. Note that the **Formula Editor** dialog displays when you click the ellipsis (**...**) button to the right of the **Value** text box in the **Values** pane. See HTTP Content Type Volumes layout for an example of this configuration.

  The **Formulas** that are provided in Message Analyzer by default for operating upon field values are described in the **Formula** list below. Immediately following are descriptions of the controls in the **Values** pane of the **Edit Chart Layout** dialog:

  - **Label** — a text box in which you to specify a name for the **Values** of fields that will display in the **Layout**. Currently, this feature is most useful with the **Table** grid visualizer component, where the **Label** value that you specify for each field or formula added to the **Values** pane displays as a table column name in your **Layout**. However, you can utilize **Label** values to display specified data names in any visualizer component.
  - **Value** — a read-only text box that displays an operation name along with a field to which a specified formula or calculation will be applied, depending on the configuration that you create in the **Formula Editor**. For example, an operation such as **Cumulative Addition** upon a **PayloadLength** field displays in the **Value** text box of the **Values** pane as **Sum(PayloadLength)**. Note that it also displays in a label in the lower sector of the **Formula Editor** dialog.
  - **Ellipsis** — a button that enables you to display the **Formula Editor** dialog, from where you can

specify a formula that operates upon an **Argument type** that consists of a message field, computed value, or a constant; along with an **Argument value** that could be a field, configured formula, or constant value that you specify with **Field Chooser**, another **Formula Editor** instance, or by manual entry, respectively. When you click the ellipsis button, the **Formula Editor** dialog appears with the following controls that enable you to perform the indicated tasks:

- **Formula** — a drop-down list that enables you to select one of seven different operations to perform on one or more fields, for example, **Count of Occurrences** for a single field and **Subtraction** or **Division** for two fields. The operations that you can specify in a **Formula** consist of the following:
  - **Average** — enables your visualizer component to display an average value for particular field, property, or annotation.
  - **Count of Occurrences** — enables your visualizer component to display one or more elements that represent the count value for a particular type of field, property, or annotation, for example, module count.
  - **Subtraction** — enables your visualizer component to calculate and display the difference between two field or property values such as a start time and end time.
  - **Division** — enables your visualizer component to perform a division operation on two field or property values, for example, to divide some value by a factor of 1000 with a constant.
  - **Maximum** — enables your visualizer component to display the maximum value of a particular field, property, or annotation.
  - **Minimum** — enables your visualizer component to display the minimum value of a particular field, property, or annotation.
  - **Cumulative Addition** — enables your visualizer component to specify the sum of values for a particular field, property, or annotation.
- **Argument** pane — contains the following controls:
  - **Argument type** — a drop-down list that contains the following three items for selection:
    - **Message Field** — choose this item when the entity to which your formula will apply is a message field that inherently displays field values.
    - **Computed Value** — choose this item when you want to present a calculated value based upon an operation that manipulates the values of one or more fields.
    - **Constant** — choose this item when you need to create a constant value for use in other operations.
  - **Argument value** — provides an ellipsis (**...**) button that displays either the **Field Chooser** if you specified the **Message Field** item in the **Argument type** drop-down list, or another instance of the **Formula Editor**, if you specified the **Computed Value** item in the **Argument type** drop-down list.

    In the former case, you will typically be able to specify a message field from the **Field Chooser** and an applied operation such as **Average**, **Count of Occurrences**, or **Cumulative addition**.

    In the latter case, you will typically be able to specify a computed value based on two

Argument values along with an operation such as **Division**, **Subtraction**, or possibly **Minimum** and **Maximum** values.

> **NOTE**
> A label displays below the **Argument** pane of the **Formula Editor** dialog that indicates the operation that you specified and the field/s upon which the operation will act, for example, **Sum(PayloadLength)**. The text of this label also appears in the **Value** text box in the **Values** pane of the **Edit Chart Layout** dialog.

  - **Add Field** — this button displays only when you select the **Table** grid visualizer component in the **Chart type** drop-down list in the **Chart Properties** pane of the **Edit Chart Layout** dialog. Enables you to display additional sets of message field controls when clicked. Each message field that you specify becomes a new column in the **Table** grid component under a **Label** name that you specify.
  - **Add Value** — this button displays only when you select the **Table** grid visualizer component in the **Chart type** drop-down list in the **Chart Properties** pane of the **Edit Chart Layout** dialog. Enables you to display additional sets of **Value** controls when clicked. Each **Value** that you create with a formula becomes a new column in the **Table** grid component under a **Label** name that you specify.

- **Apply** button — after you complete your **Layout** configuration with the **Formula Editor** and **Edit Chart Layout** dialogs, click the **Apply** button to render the results of your configuration as a new **Chart** viewer **Layout**. If the results meet your expectations, you can save your customized **Chart** as specified in Saving a Custom Chart Viewer Layout. Otherwise, you can return to editing by again displaying **Edit Chart Layout** dialog or you can close the **Chart** viewer **Layout** without saving your changes.

Caution

If you close the **Chart** viewer tab on which your modified **Layout** exists, you will lose the **Layout** configuration and you will be unable to recover it unless you reconfigure the **Layout**.

Saving a Custom Chart Viewer Layout

To save your customized **Chart** viewer **Layout**, click the **Layout** command in the **Charts** drop-down list in the global Message Analyzer **Session** menu and then select the **Save Current Layout As...** command to display the **Edit Chart Layout** dialog that has your custom configuration along with **Name**, **Description**, and **Category** fields so that you can rename, describe, and categorize your customized **Layout** before saving it. After you provide this information, click the **Save** button to save your **Layout** and exit the **Edit Chart Layout** dialog.

> **NOTE**
> Because Message Analyzer does not enable you to overwrite any of the built-in **Chart** viewer **Layouts**, you can save any modifications you have made to a built-in **Layout** as a new **Layout** only.

Note that if you modify and save a built-in **Chart** viewer **Layout** under a specified name, it will be saved in the **My Items** category of the **Message Analyzer Charts** asset collection that you access from the categorized **Layout** drop-down list in the global Message Analyzer **Session** menu. Your new **Layout** will also appear in the uncategorized lists that display in the following locations:

- The **New Viewer** drop-down list on the global Message Analyzer toolbar.
- The **Chart** drop-down list in the **New Viewer** drop-down of the **Session Explorer** context menu.

Caution

If you elect to not save your custom **Layout** configuration and you close the **Chart** viewer, you will lose all the configuration settings that you specified.

See Also

Configuration Walkthrough of a Built-In Chart Viewer Layout

Configuration Walkthrough of a Built-In Chart Viewer Layout

15 minutes to read

This section provides a walkthrough of the configuration process for the built-in **Chart** viewer **Layout** known as the TCP/UDP Conversations by Message Count**Layout**. You might click the specified link and review the indicated topic in order to understand the functionality of this **Layout** before proceeding with the walkthrough.

Built-In Chart Viewer Layout Configuration Walkthrough

The **TCP/UDP Conversations by Message Count Layout** is accessible from the **Chart** drop-down list in the **New Viewer** menu on the global Message Analyzer toolbar. This **Layout** enables you to view the network conversations and the transports that carried them, along with the following data that can point you to specific areas that may require further investigation:

- Conversations with the highest message volume
- Conversations with the highest payload volume
- Conversations with the highest data transmission rates
- Conversations with the highest durations

The data that you evaluate to facilitate this analysis appears in the columns of the **TCP/UDP Conversations by Message Count Layout** in tabular format, with the use of the **Table** grid visualizer component. The data consists of the following:

- **Network conversation** data — displays in the **Network** column of the Table and exposes the IPv4, IPv6, and/or Ethernet addresses of the computers engaged in each conversation across a set of trace results.
- **TCP/UDP transport** data — displays in the **Transport** column of the Table and exposes the ports on each computer that carried the conversations.
- **Message count** data — displays in the **Count** column of the Table and exposes the cumulative message count that is calculated for each conversation.
- **Payload statistics** data — displays in the **Bytes** column and exposes the calculated sum of the total number of bytes transmitted in each conversation.
- **Data transmission rate** data — displays in the **KBS** and **BPS** columns of the Table and provides a computed data transmission rate in kilobytes-per-second and bytes-per-second, respectively, for each conversation.
- **Duration statistics** data — displays in the **Duration** column and provides a computed time interval for the duration of each conversation.

Configuration Walkthrough

The subsections below explain how the fields and computed values that populate the Table columns of the **TCP/UDP Conversations by Message Count Layout** are configured with the use of the **Edit Chart Layout** and **Formula Editor** dialogs. To get the most out of this walkthrough, you should display this **Layout** and open the **Edit Chart Layout** and **Formula Editor** dialogs in the procedures and explanations that follow, so that you can follow the process step-by-step through the dialog configurations that are given here.

To start Message Analyzer, load data, and display the TCP/UDP Conversations by Message Count Layout

1. Start Message Analyzer and open a saved *.matp file that contains data that you recently captured, preferably with a **Trace Scenario** that uses the **Microsoft-PEF-WFP-MessageProvider** to minimize lower layer noise, given that the main analysis will be at the Transport Layer with this **Layout**.

   You can load the data from a *.matp file into Message Analyzer through a Data Retrieval Session or with the **Open** command on the global Message Analyzer toolbar.

2. After the data is loaded, click the **New Viewer** drop-down list on the global Message Analyzer toolbar, highlight **Chart**, and then click **TCP/UDP Conversations by Message Count**.

   The **TCP/UDP Conversations by Message Count Layout** should display with data populated in the **Table** grid visualizer component.

3. While the **TCP/UDP Conversations by Message Count Layout** has focus, click the global Message Analyzer **Session** menu, highlight **Chart**, and then click **Edit** in the drop-down list to display the **Edit Chart Layout** dialog.

Chart Properties Configuration

In the **Chart Properties** pane of the **Edit Chart Layout** dialog, you will see that the **Chart type** is set to **Table**, the **Sort Order** is set to **Descending**, and **Sort By** is set to **Values**. This means that quantities in the first sortable **Value** column, in this case the **Count** column of this **Layout**, will be sorted in descending order by default, so that you can assess the network conversations from the highest message counts and payloads to the lowest, as you scroll down through the data.

The data that displays in the columns of the Table visualizer component for this **Layout** are described in the subsections that follow.

Network Conversations

For this **Layout** to display data in the **Network** column of the Table, a **Field** entry in the **Series Field** pane must be configured with the **Network** field from the IPv4 or IPv6 protocol that you can locate in the **Field Chooser**. Note that the **Network** field, like all other fields in **Field Chooser**, contains an *inherent value* that Message Analyzer can display in various data viewers where it is used, including in **Chart** viewer **Layouts**. This contrasts with a *calculated value* that is the result of a selected operation, such as **Cumulative addition**, upon one or more data fields that produces a statistic that is useful for analysis. Adding fields with inherent values is the only type of configuration that is allowed in the **Series Field** pane, given that the **Formula Editor** is unavailable in this context.

To locate the Network field

1. In the **Series Fields** pane of the **Edit Chart Layout** dialog, click the ellipsis button next to the first **Field** text box to open the **Field Chooser** window.

2. In the **Field Chooser** window, scroll down to the **IPv4** or **IPv6** node and then click the expansion control to display the **Datagram** message type.

   If the **Field Chooser** is not already displayed, you can find it in the **Windows** drop-down list that is accessible from the global Message Analyzer **Tools** menu.

3. Click the expansion control of the **Datagram** node to expose the **IPv4** message field hierarchy.

4. Scroll down to the **Network** field in the hierarchy.

   If this was an actual configuration task rather than a walkthrough, you would need to double-click the **Network** field that you located in **Field Chooser** to display the "Network" value in the first **Field** text box of the **Series Fields** pane in the **Edit Chart Layout** dialog.

TCP/UDP Transports

For this **Layout** to display data in the **Transport** column of the Table, a second **Field** entry in the **Series Fields** pane must be configured with the **Transport** field from the TCP or UDP protocol, which you can locate in the

Field Chooser window.

To locate the Transport field

1. In the **Series Fields** pane of the **Edit Chart Layout** dialog, click the ellipsis button next to the second **Field** text box to open the **Field Chooser** window.
2. In the **Field Chooser**, scroll down to the **TCP** protocol and then click the expansion control to display the **Segment** message type.
3. Click the expansion control of the **Segment** node to expose the **TCP** message field hierarchy.
4. Scroll down to the **Transport** field in the **TCP** hierarchy.

   If this was an actual configuration task rather than a walkthrough, you would need to double-click the **Transport** field that you located in **Field Chooser** to display the "Transport" value in the second **Field** text box of the **Series Fields** pane in the **Edit Chart Layout** dialog.

Message Count

For this **Layout** to display data in the **Count** column, the first **Value** field in the **Values** pane must be configured with a **Count of Occurrences** operation in the **Formula Editor**. The formula for this **Value** field sets the **Argument type** to the **Message Field** option and sets the **Argument value** to **MessageNumber**, the latter of which you locate in **Field Chooser**. The desired outcome for this configuration is to provide the total message count that is associated with each TCP or UDP conversation.

Note that **MessageNumber** is a **Global Annotation** in **Field Chooser** that you can use to count messages, given that Message Analyzer assigns a **MessageNumber** to each message it parses. Therefore, in this **Layout**, the count of **MessageNumbers** is equivalent to the number of messages for a particular conversation.

To compute message Count values

1. In the **Values** pane of the **Edit Chart Layout** dialog, click the ellipsis button next to the first **Value** text box to open the **Formula Editor** dialog.

   Note the following settings in the **Formula Editor** dialog:

   - **Formula** — set to **Count of Occurrences**, which enables the **Layout** to count the occurrences of a particular entity or value; in this case it will be the message **Count**.
   - **Argument type** — an **Argument** pane option that is set to **Message Field**, which enables the **Layout** to count the occurrences of a particular field that is selected in **Field Chooser**, as specified below.
   - **Argument value** — an **Argument** pane option that is set to **MessageNumber**; this entity is located under the **Global Annotation** node in **Field Chooser**, which displays when you click the ellipsis next to the **Argument value** text box. As stated earlier, by using the **MessageNumber** annotation and the **Count of Occurrences** operation in a formula, you can count the number of messages and display the result in any visualizer component.
2. Observe that the **Formula Editor** dialog contains a label below the configuration controls that specifies the resulting operation that will be performed on a specified entity, for example, **Count(MessageNumber)** in the case of the above configuration. This formula also appears in the first **Value** text box in the **Values** pane after you click **OK** to exit the **Formula Editor** dialog.

Payload Statistics

For this **Layout** to display data in the **Bytes** column, the second **Value** field in the **Values** pane must be configured with a **Cumulative addition** operation in the **Formula Editor**. The formula for this **Value** field sets the **Argument type** to the **Message Field** option and sets the **Argument value** to **PayloadLength**, which you locate in **Field Chooser**. The desired outcome for this configuration is to provide the cumulative sum of payload lengths in bytes for each TCP or UDP conversation, from all messages that define a **Payload** field.

Note that although there are protocols that do not have a **Payload** field, in which case **PayloadLength** could not be evaluated, the TCP and UDP protocols do have such a field.

To compute the sum of message payloads in Bytes

1. In the **Values** pane of the **Edit Chart Layout** dialog, click the ellipsis button next to the second **Value** text box to open the **Formula Editor** dialog.

   Note the following settings in the **Formula Editor** dialog:

   - **Formula** — set to **Cumulative addition**, which enables this **Layout** to create the sum of all instances of a particular value; which in this case is the cumulative message **PayloadLength** in bytes for each conversation.
   - **Argument type** — an **Argument** pane option that is set to **Message Field**, which enables this **Layout** to create the sum of all instances of a particular field value that is specified in the **Argument value** text box. You can select the field for which you want to sum all values in **Field Chooser**, as specified below.
   - **Argument value** — an **Argument** pane option that is set to **PayloadLength**; this entity is located in the hierarchy of many message types in **Field Chooser** that define a **Payload** field, for example, HTTP, TCP, or IPv4. **Field Chooser** displays when you click the ellipsis next to the **Argument value** text box.

2. Observe that a **Formula Editor** dialog label below the configuration controls specifies the resulting operation that will be performed on the specified field, for example, **Sum(PayloadLength)** in the case of the above configuration. This formula also appears in the second **Value** text box in the **Values** pane after you click **OK** to exit the **Formula Editor** dialog.

Conversation Data Transmission Rate

For this **Layout** to display data in the **KBs** column, the third **Value** field in the **Values** pane must be configured by several layers of operations with multiple instances of the **Formula Editor**. The desired outcome for this configuration is to generate the rate in kilobytes, at which data was transmitted in a conversation. A word description of the final formula that implements the calculations consists of the following parts:

- A **Subtraction** operation that computes the difference between the conversation **StartTime** and **EndTime** time stamps.
- A **Division** operation that divides the sum of **PayloadLengths** by the value that results from the above **Subtraction** operation.
- A final **Division** operation that factors the result of the previous **Division** operation by 1000, to obtain a kilobytes-per-second (KBS) value for display in the **KBs** column of this **Layout** for each conversation.

To obtain the results of these operations, multiple formulas are required, as described in the procedure that follows.

To compute conversation data transmission rates in kilobytes (KBs)

1. In the **Values** pane of the **Edit Chart Layout** dialog, click the ellipsis button next to the third **Value** text box to open the **Formula Editor** dialog, as shown in the figure that follows.

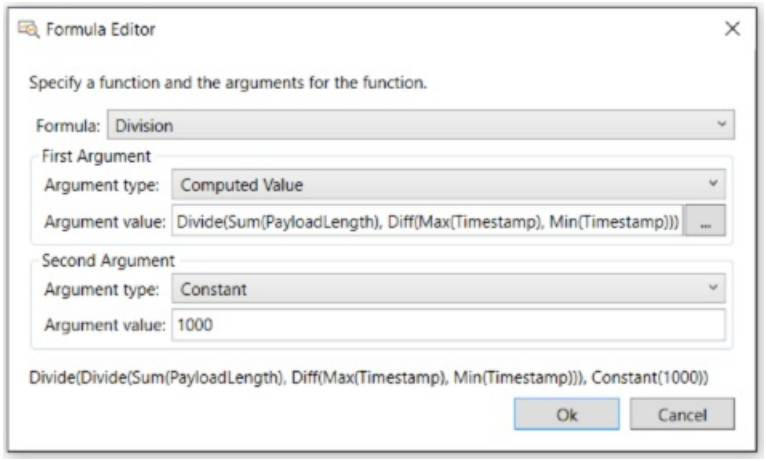

Figure 73: Formula Editor dialog

Note the following settings in this dialog:

- The **Formula** drop-down list is set to the **Division** operation.

  Note that when either a **Division** or **Subtraction** operation is selected in the **Formula Editor** dialog, two sets of **Argument** controls display, one for the **First Argument** and one for a **Second Argument**.

- The **Argument type** drop-down list in the **Second Argument** pane is set to the **Constant** item and the **Argument value** text box is manually configured with a constant value of 1000, which is needed to convert the **PayloadLength** sum in bytes-per-second over the conversation duration into a *kilobytes-per-second* value.

- The **Argument type** drop-down list in the **First Argument** pane is set to the **Computed Value** item and the **Argument value** text box displays a formula that is created through multiple **Formula Editor** dialog configurations that are described in the next few steps.

2. In the first **Formula Editor** dialog displayed thus far, click the ellipsis button next to the **Argument** value text box in the **First Argument** pane to display another **Formula Editor** dialog instance that is set to perform a **Division** operation using the following two **Computed Values**:

   - **Sum(PayloadLength)** — the configuration for creating the sum of **PayloadLength** values in bytes for each conversation is described earlier in the **Payload Statistics** subsection. However, you can look at the **Formula Editor** for this configuration again by clicking the ellipsis button next to the **Argument value** text box in the **First Argument** pane of the current **Formula Editor** dialog instance.

     The **Formula Editor** dialog that displays is configured with a **Division** operation on two **Computed Values**, where the **Sum(PayloadLength)** value will be divided by the **Diff(Max(Timestamp), Min(Timestamp))** value, which is described in the next bullet point.

   - **Diff(Max(Timestamp), Min(Timestamp))** — the configuration for subtracting the difference between the **StartTime** and **EndTime** time stamps, which is needed to compute the **PayloadLength** per-second data transmission rate for each conversation, is exposed in the **Formula Editor** dialog instance that displays when you click the ellipsis button next to the **Argument value** text box in the **Second Argument** pane of the current **Formula Editor** dialog instance.

3. Observe in the current **Formula Editor** dialog instance (from the last bullet point) that a **Subtraction** operation is set to subtract two **Computed Values** that are configured as **Max(Timestamp)** and **Min(Timestamp)**. You can view the configuration for these formulas that use either a **Maximum** or **Minimum** operation, respectively, by clicking the ellipsis button next to the formula in each text box.

Conversation duration

For this **Layout** to display data in the **Duration** column, the fourth **Value** field in the **Values** pane must be configured by several layers of operations with multiple instances of the **Formula Editor**. The desired outcome for this configuration is to provide a statistic that exposes how long the message exchanges in each conversation took to complete. A word description of the final formula that implements the calculations consists of the following parts:

- A **Subtraction** operation that computes the difference between the maximum time stamp value and the minimum time stamp value in a specific conversation.
- A **Maximum** operation that computes a result that is equal to the maximum time stamp value in the conversation.
- A **Minimum** operation that computes a result that is equal to the minimum time stamp in the conversation.

To obtain the results of these operations, multiple formulas are required, as described in the procedure that follows.

To compute conversation Durations

1. In the **Values** pane of the **Edit Chart Layout** dialog, click the ellipsis button next to the fourth **Value** text box to open the **Formula Editor** dialog.

   Note the following settings in the **Formula Editor** dialog:

   - The **Formula** drop-down list is set to the **Subtraction** operation.
   - The **Argument type** drop-down list in the **First Argument** pane is set to **Computed Value** and the **Argument value** text box is set to the formula **Max(Timestamp)**.
   - The **Argument type** drop-down list in the **Second Argument** pane is also set to **Computed Value** and the **Argument value** text box is set to the formula **Min(Timestamp)**.

2. To review how the formula **Max(Timestamp)** or **Min(Timestamp)** is configured, click the ellipsis button next to the **Argument value** text box in the **First Argument** or **Second Argument** pane, respectively.

3. If you click the ellipsis button in the **First Argument** pane of the step 2, observe in the current **Formula Editor** dialog instance that a **Maximum** operation is set to calculate the value of the latest time stamp in a conversation.

   If you now click the ellipsis button next to the **Argument value** of the current **Formula Editor** dialog instance, you will display the **Field Chooser**, which provides for the selection of the **TimeStamp** annotation for the **Maximum** value.

4. If you click the ellipsis button in the **Second Argument** pane of step 2, observe in the current **Formula Editor** dialog instance that a **Minimum** operation is set to calculate the value of the earliest time stamp in a conversation.

   If you now click the ellipsis button next to the **Argument value** of the current **Formula Editor** dialog instance, you will display the **Field Chooser**, which provides for the selection of the **TimeStamp** annotation for the **Maximum** value.

> **NOTE**
>
> There are other **Formulas** that you can review for data that displays in the **StartTime**, **EndTime**, **BPS**, and **K** columns of this **Layout**. However, an explanation of the **Formulas** for this data has already been provided in other parts of this walkthrough.

See Also

Using the Edit Chart Layout Dialog

Managing Chart Viewer Layouts

7 minutes to read

This section describes how to manage the items of your **Message Analyzer Chart View Layouts** asset collection. Also included are discussions about using the Message Analyzer Sharing Infrastructure to share **Chart** viewer **Layouts** with others and how to receive automatic updates for this collection.

Managing the Chart Viewer Layouts Library

To manage the **Chart** viewer **Layouts** in the **Message Analyzer Chart View Layouts** asset collection Library, you will use commands that are available in the **Manage Chart Layout** dialog, which is accessible from the global Message Analyzer **Session** menu by selecting the **Chart**, **Layout**, **Manage Layouts**, and **Manage** items in succession.

The **Message Analyzer Chart View Layouts** asset collection Library has a **Message Analyzer** category for the default asset collection that contains the built-in **Layouts** that are described in the Built-In Chart Viewer Layouts topic. This Library also contains a **My Items** category from where you can manage any **Layouts** of your own design that you saved. The commands that are available in this dialog are described in Managing Chart Viewer Layout Library Items.

All the Library collection items in both of these categories are shareable **Layout** items. Message Analyzer provides a simple way to expose your **Layouts** asset collection items to others for sharing or to retrieve **Layouts** that others have shared. You can share **Layout** collection items directly with others by using the **Export** feature in the **Manage Chart Layout** dialog to save one or more **Layout** collection items to a designated file share. You can also use the **Import** feature in the same dialog to access **Layout** collection items that have been shared by others.

More Information
To learn more about managing Message Analyzer asset collections, including the **Message Analyzer Chart View Layouts** collection, see Managing User Libraries.

Using the Layout Commands

This section briefly describes the Layout commands that you can use to manage any of the built-in **Chart** viewer **Layouts** or any new **Layout** that you have created. These commands are accessible from the **Layouts** drop-down list that appears on the Charts toolbar whenever a **Chart** viewer **Layout** is displayed. A descriptions of each command is specified in the list that follows:

- **Save Current Layout As...** — the primary command that you will use to save a newly modified **Chart** viewer **Layout** as a custom **Layout** of your own. When you click this command from the **Layout** drop-down list on the Charts toolbar, the **Edit Chart Layout** dialog displays with the **Layout** configuration and formulas that you specified, where you can provide a **Name**, **Description**, and a **Category** for your custom **Chart** viewer **Layout**.
- **Load Default User Layout** — appears in the **Manage Layouts** submenu of the **Layout** drop-down list that displays on the Charts toolbar. Enables you to load the default user **Layout** that you previously specified with the **Save Current as Default User Layout** command.
- **Restore Application Default Layout** — appears in the **Manage Layouts** submenu of the **Layout** drop-down list that displays on the Charts toolbar. Enables you to load the application default **Chart** viewer **Layout**, which consists of a bar element visualizer that displays the relative distribution of message volume

for each protocol or module that was captured in a set of trace results.

- **Save Current as Default User Layout** — appears in the **Manage Layouts** submenu of the **Layout** drop-down list that displays on the Charts toolbar. Enables you to save the currently displayed **Chart** viewer **Layout** as the default, such that you can display it any time thereafter by executing the **Load Default User Layout** command.

Utilizing the Message Analyzer Sharing Infrastructure

You can share your **Layout** collection items through the Message Analyzer Sharing Infrastructure by using the **Export** feature of the **Manage Chart Layout** dialog to post one or more **Layouts** to a file share or other designated location that is accessed through a user-configured feed. You can create your own feed from the **Settings** tab in the **Asset Manager** dialog, which is accessible from the global Message Analyzer **Tools** menu. Thereafter, you can update existing **Layout** collection items or add others to make them available to team members or other users through the configured feed, where they can view, synchronize, and download them. However, the synchronization feature that keeps users up to date requires some manual configuration at this time, as described in Manual Item Update Synchronization.

Message Analyzer also has a default subscriber feed that appears on both the **Downloads** and **Settings** tabs of the **Asset Manager** dialog. On the **Downloads** tab, the feed enables you to view the default assets provided with Message Analyzer. On the **Settings** tab, the feed provides the path that enables the Message Analyzer application to download the **Message Analyzer Chart View Layouts** asset collection from a Microsoft web service at first Message Analyzer startup. It also enables you to receive asset collection updates that are periodically pushed out by the service, as useful **Layout** assets are developed at Microsoft for the community of Message Analyzer users. To receive updates that will appear in the **Message Analyzer** category of your local **Message Analyzer Chart View Layouts** asset collection Library, you must set this asset collection to the auto-sync state on the **Downloads** tab of the **Asset Manager**, as described in Downloading Assets and Auto-Syncing Updates.

Managing Chart Viewer Layout Library Items

You can manage **Layout** items in both categories of your local **Message Analyzer Chart View Layouts** Library. The following features are available to do so:

- **Item selection** — you can select **Layout** collection items to include in an export configuration by selecting all items in the **Message Analyzer** category, the **My Items** category, or both. You can also select any combination of individual collection items in these categories to include in the export.
- **Context menu commands** — the following context menu command is available by right-clicking any collection item in the **Message Analyzer** category of your local **Layouts** asset collection Library:
  - **Create a Copy** — enables you to create a copy of an existing **Layout** collection item, such as the **Average Response Times for Operations** layout, and place it in your **My Items** category. You can then launch this copied **Layout** in an Analysis Session and use the controls and features of the **Edit Chart Layout** and **Formula Editor** dialogs to edit and customize the **Layout** to your own design and then save it under a new name of your choice, as described in Overview of Configuring a Custom Chart Viewer Layout.

The following context menu commands are available by right-clicking any collection item in the **My Items** category of your local **Layouts** asset collection Library:

```
-   **Edit** – enables you to save a **Layout** collection item with a new **Name**, **Description**, and
**Category**.

-   **Create a Copy** – enables you to create a copy of a **Layout** collection item and use it as  a
functionality template for creating a new **Layout** that will have similar components in the configuration.

-   **Delete** – enables you to delete any **Layout** collection item in the **My Items** category. When you
select this command, all selected **Layout** collection items are immediately deleted without any prompt. Note
that you cannot delete any Library collection items in the **Message Analyzer** category.
```

- **Import** — enables you to open the **Select Library to Open** dialog from where you can navigate to a **Layout** asset collection. When you exit this dialog by clicking the **Open** button, the **Select Items to Import** dialog opens, from where you can choose the **Layout** items you want to import and the category into which they will be placed. Note that you can navigate to a user-designated file share or other location from the **Select Library to Open** dialog, to import a **Layout** asset collection that another user has posted in that location.

- **Export** — enables you to open the **Save Library** dialog, from where you can specify a library **Title** along with optional **Description**, **Author**, and **Organization** information. When you exit this dialog by clicking the **Save** button, the **Select Library Location** dialog opens, from where you can navigate to a user-designated file share or other location to post your **Layout** asset collection files.

More Information

To learn more about sharing Message Analyzer asset collection Library items, including further details about the common **Manage <*AssetType*>** dialog, see the Sharing Infrastructure topic.

To learn more about the manual synchronization process for a user-configured feed, see Manual Item Update Synchronization.

To learn more about auto-syncing and downloading item collections, see Managing Asset Collection Downloads and Updates.

Procedures: Using the Chart Configuration Features

4 minutes to read

This section provides an example of how to extend Message Analyzer data viewing capabilities by creating a custom **Chart** data viewer that you can use to analyze message volume against HTTP payload content types.

Procedure Overviews
A brief description of the procedure in this section is included below for review. Note that more **Chart** development procedures will be provided in the future.

Create and Edit an HTTP Chart Data Viewer — demonstrates how you can create your own chart-style data viewer for enhanced data analysis perspectives. This particular **Chart** exposes the message volume associated with various HTTP content types being delivered to a client computer in HTTP response messages, which can provide an indication of the load on a web

Create and Edit an HTTP Chart Data Viewer

This example demonstrates how to create a new **Chart** viewer that is named "HTTP Content Types", which provides a **Bar Chart** visualizer component that shows the relative volume distributions of the content types for HTTP messages captured in a trace.

To create and edit a Chart viewer

1. From the **Start** menu, **Start** page, or taskbar of your computer, click the **Microsoft Message Analyzer** icon to start Message Analyzer.
2. From **Favorites** list on the Message Analyzer **Start Page**, click the **Pre-Encryption for HTTPS Trace Scenario** and then navigate to one or more web sites to generate browser traffic and capture HTTP messages.

   Message Analyzer begins to collect HTTP and other messages in the **Analysis Grid** viewer.
3. At a suitable point, stop the trace by clicking the **Stop** button on the global Message Analyzer toolbar.
4. From the **New Viewer** drop-down list on the global Message Analyzer toolbar, click the **New Chart** item to display the **Chart** visualizer surface and data entry tabs.
5. In the **Name** property text box on the **Chart Layout** tab to the right of the **Chart** visualizer surface, enter the text "HTTP Content Types" to specify the name that will display for your new **Chart** viewer in the **New Viewer** drop-down list.
6. In the **Component Common** section of the **Component Layout** tab to the right of the **Chart** visualizer surface, specify a title for your **Chart** viewer in the **Name** text box.
7. In the **Component Common** section of the **Component Layout** tab to the right of the **Chart** visualizer surface, click the **Type** drop-down list and select the **Bar Chart** visualizer component.
8. On the **Data** tab to the right of the **Chart** visualizer surface, click the ellipsis (**...**) on the right side of the **Column Fields** property box to display the **Data Mapping Field Collection Editor** dialog.
9. In the **Data Mapping Field Collection Editor** dialog, click the **Add** drop-down list and select the **Formula** menu item.

   A new **Formula** item appears in the **Members** pane.

10. Click the **FormulaType** drop-down list in the **Formula properties** pane of the **Data Mapping Field Collection Editor** and select the **Count** item.

11. In the **Formula properties** pane of the **Data Mapping Field Collection Editor** dialog, click the ellipsis (**...**) to the right of the **Arguments** property to display the **Argument Collection Editor**.

12. In the **Argument Collection Editor**, click the **Add** button to add an **Argument** item to the **Members** pane.

13. In the **Argument properties** pane of the **Argument Collection Editor**, click the **FieldName** ellipsis (**...**) to display the **Field Chooser Tool Window**.

14. In the **Field Chooser** window, scroll down to the **HTTP** node, expand it, and select the **HTTP.Operation.ContentType** field.

15. In the **Field Chooser**, click the **Select** button to exit the dialog.

    The **ContentType** field displays in the **Arguments** list of the **Data Mapping Field Collection Editor**.

16. Click **OK** to exit the **Data Mapping Field Collection Editor**.

17. On the **Data** tab, click the **Row Fields** ellipsis (**...**) to display the **Data Mapping Field Collection Editor** again, click the **Add** button to display the **Entity** item. A new **Entity** displays in the **Members** pane of the editor.

18. In the **Entity properties** pane, click the **Name** ellipsis (**...**) to display the **Field Chooser** window.

19. In the **Field Chooser** window, scroll down to the **HTTP** node, expand it, and select the **HTTP.Operation.ContentType** field.

20. In the **Field Chooser** window, click the **Select** button to exit the dialog.

    The **ContentType** field displays as the **Name** property in the **ContentType properties** pane of the **Data Mapping Field Collection Editor**.

21. Click **OK** to exit the **Data Mapping Field Collection Editor**.

    Your new **HTTP Content Types Chart** displays each different HTTP content type as a separate horizontal bar in the **Bar Chart** visualizer component, with each bar representing the count of HTTP messages that carried the indicated **ContentType** payload, for example, image/gif, text/html, image/jpeg, and so on.

    To view the messages that correspond with the different content types, double-click any bar and the messages display in a separate **Analysis Grid** viewer tab.

22. Save your existing **Chart** configuration by clicking the **Save Chart** command in the **Charts** submenu that appears in the global Message Analyzer **Session** menu.

    Your new **Chart** appears in all of the locations described in Built-In Chart Data Viewers, from where you can select the **Chart** as required for simple HTTP analysis.

23. To edit your new **Chart** at any time, select it from the **New Viewer** drop-down list against a set of trace results and then click the **Edit Chart** item in the **Charts** submenu that appears in the global Message Analyzer **Session** menu.

    Your **Chart** enters the edit mode and the data entry tabs appear to the right of the **Chart** visualizer surface. From here, you can modify your **Chart** settings with the use of the **Add Component**, **Remove Component**, and **New Data Mapping** commands from the **Charts** submenu that appears in the global **Session** menu. Note that you can also modify settings directly from the data entry tabs.

Participating in the Message Analyzer Community

2 minutes to read

As a Message Analyzer user, you can take advantage of community resources to enhance your knowledge of Message Analyzer and network troubleshooting. To this end, you can participate in Message Analyzer Forum discussions to obtain feedback on issues or to report bugs. You can also respond to Message Analyzer team Blog postings. These features enable you to connect with Microsoft and other users so you can benefit from their perspectives and experiences with Message Analyzer. There are also facilities in the Message Analyzer user interface to provide feedback about product features. In addition, you can provide feedback on specific topics of this Operating Guide in either of the following ways, while in the case of the first bullet point below, you can receive feedback directly from Microsoft:

- Use the **Community Additions** feature at the end of each topic in this Operating Guide to start a comment and response thread, where you can receive direct feedback from Microsoft.
- Use the **Was this page helpful?** feature at the end of each topic in the Operating Guide to provide your comments, although Microsoft cannot directly respond to comments when you use this option.

What You Will Learn

In the topics of this section, a brief description of the above specified resources is included. Enhancements to the way you can provide **Feedback** directly from the Message Analyzer user interface are also described.

In This Section

Message Analyzer Community Additions — learn how you can provide feedback or start a discussion thread on any topic in this Operating Guide, to which you can obtain feedback directly from Microsoft.

Message Analyzer Team Blog — take advantage of the Message Analyzer Blog site, where informative articles are posted by the engineers who designed Message Analyzer.

Message Analyzer Online Forum — create a discussion thread to obtain feedback from Microsoft on bug related issues and others; also enables you to facilitate problem solving by soliciting interactive participation from other users on your issue.

Message Analyzer Feedback — review all the options that are available for providing feedback about Message Analyzer, including on features you have used and features that you want to request.

Message Analyzer Community Additions

2 minutes to read

At the end of each topic in the Message Analyzer Operating Guide, there is a **Community Additions** section where you can start a comment and response thread. If you provide a comment, question, or legitimate concern that is not spam, you will receive direct feedback from Microsoft. To get started, you simply click the **ADD** link in the **Community Additions** section. Note that you will be prompted to specify Microsoft account credentials before you can proceed. Thereafter, a **Community Additions** page displays with text boxes that enable you to specify a title and the content for your community entry. When complete, you simply click the **Submit** button and your entry will appear in the **Community Additions** section at the end of the topic for which you provided comments.

Message Analyzer Team Blog

2 minutes to read

Periodically, members of the Microsoft PEF engineering team post interesting articles on the Message Analyzer Team Blog on TechNet. From the blog site, you can learn about Message Analyzer features from the perspective of those who designed it, which can be invaluable to your Message Analyzer expertise. You will also have the opportunity to ask questions and leave comments about blog posts and you can interact with other Message Analyzer users who may have already solved an issue that you are dealing with. Also, the site is monitored by the PEF engineering team, so you can be assured that your questions and concerns will be addressed.

Go To the Message Analyzer Team Blog

Blog Feature Summary

The following Message Analyzer blog feature summary describes some of the specific things you can do at the site:

- **Review blog posts** — read articles that provide tips, tricks, and how-to information for using Message Analyzer features.
- **Rate blog posts** — click one of the five stars above the article to specify a rating of 1 to 5, with 5 being the highest rating.
- **Leave a comment** — post a comment or question against a blog post.
- **Sort blog postings** — organize your blog display by the most recent, the most viewed, or the most comments, or inspect a monthly archive category.
- **Search for a blog by name** — enter search text and search for a blog title match.
- **Share a blog post** — share a post with another Message Analyzer user by sending them a blog title and link via email, or share a post with others on Facebook, Twitter, Digg, LinkedIn, and other social media sites, by sending them a blog title and link.
- **Subscribe to feeds** — options include RSS for posts, Atom, and RSS for comments:
  - Subscribe to the blog posts feed, so that you can receive information updates about new blog postings that are automatically downloaded to your computer.
  - Subscribe to the blog comments feed, so that you can receive questions/comments from other users and engage in the discussions on posted blog articles.

Blog Posts

A few sample blog postings describing Message Analyzer features are included here so you can get started:

Diagnosing Network Issues with the New Pattern Match Viewer
Message Analyzer v1.3 vs Network Monitor v3.4
Relating and Correlating Data Sources
Process Tracking with Message Analyzer
Parsing Text Logs with Message Analyzer Grouping Viewer - AKA Conversation Tree

NOTE

Many more blog posts will be forthcoming as the community of Message Analyzer users continues to ramp up.

Message Analyzer Online Forum

2 minutes to read

If you have comments or questions about Message Analyzer issues that you encounter, you can start a discussion thread on the Message Analyzer Forum on TechNet. This includes issues that you may experiencing such as an application crash, unknown exception, or others that appear to be bug related. Microsoft regularly monitors the Message Analyzer Forum, so if you are having an issue that is bug related or otherwise, please provide the details and Microsoft will respond.

The advantage of bringing your issue to the community of Message Analyzer users is that you can solicit interactive participation in your discussion thread to obtain feedback that may resolve a problem you are experiencing. By engaging with other users in this manner, a tribal knowledge is developed in the community that not only helps users, but enables Microsoft to make improvements to Message Analyzer based on feedback.

Go To the Message Analyzer Online Forum

Forum Feature Summary

The following Message Analyzer Forum features summary describes some of the specific things you can do at the site:

- **Start a discussion thread** — ask a question and/or provide instructions that describe the conditions under which a problem is reproduced, so that others, including Microsoft, may understand the issue and respond appropriately.
- **Edit a comment or question** — modify a comment or question that you posted in any discussion thread, which includes adding a quote, editing, or deleting your comment.
- **Participate and engage** — reply to other comment posts in a discussion thread to lend your expertise to an issue at hand.
- **Vote** — enables you to specify whether a post is helpful.
- **Search for discussions** — search for Forum questions or search on related discussion threads.
- **Filter the thread display** — specify whether to view all threads, threads that are answered, unanswered, or have proposed answers, or select the general discussion category. Additional filtering enables you to filter for threads that had no replies, those that were voted as helpful, or the ones that have code.
- **Sort the thread display** — apply sorting based on the most recent posts, the most recent threads, those with the most votes, and the ones with the most replies.
- **Alerts** — configure an email alert to show new posts for a particular discussion thread as they arrive.
- **Subscribe** — you can subscribe to a feed for a discussion thread on a particular issue, so you can automatically receive updates that are downloaded to your computer.

Message Analyzer Feedback

2 minutes to read

If you would like to share your Message Analyzer experiences with Microsoft, there are several options you can select from, as described below.

Using the Feedback List

Click the **Feedback** drop-down list that is located in the upper right section of the Message Analyzer user interface and select from the following options:

- **Send a Smile** — tell us what you liked.
- **Send a Frown** — tell us what we can do better.
- **Report a Bug** — provide us with details about problems that you encountered.
- **Request a Feature** — make a request for features that you think would improve your experiences with Message Analyzer.

Using the Feedback Center

Message Analyzer also has a **Feedback Center** that enables you to provide feedback on numerous predefined questions about your experiences with Message Analyzer. You can open the **Feedback Center** dialog by clicking the flagged yellow icon in the upper right section of the Message Analyzer user interface. To provide feedback, select a question, enter your comments in the text box, and then click the **Submit Feedback** button.

Addendum 1: Configuration Requirements for Parsing CustomText Logs

3 minutes to read

Message Analyzer automatically determines the Open Protocol Notation (OPN) configuration that is needed to parse supported input file types, which are described in Locating Supported Input Data File Types. However, in the case of textual log file types that contain proprietary message formats, it is often the case that you will need to generate an OPN configuration file that defines how your text log is to be parsed, so that you can properly display your text log data in a Message Analyzer viewer. This addendum provides a high-level overview about creating an OPN configuration file for a text log, the directory location where you must place it so that Message Analyzer can load it at startup, and how to set a configuration file to be the default for parsing all your text logs. For the low-level details required to create a configuration file, a guide is available for download, as described in the section that follows.

Creating an OPN Configuration File

An OPN configuration file consists of message definitions that describe how text log data entries are to be parsed for display in Message Analyzer. You can create the message definitions for a text log by using OPN and Regular Expression (Regex) notation to identify each unique type of log entry and map it to a message structure, as indicated in Parsing Input Text Log Files. The message definitions contained in a configuration file are also subject to OPN compilation to ensure the integrity of the OPN description that Message Analyzer will use to parse the text log. If an OPN configuration file does not properly compile, you will receive a compilation error during Message Analyzer startup when OPN definitions are loaded into the system.

You can review the details and requirements for creating an OPN configuration file by downloading the OPN Configuration Guide for Text Log Adapter document.

Saving the Configuration File

After you create a configuration file for a text log and save it with a .config extension, you must place it in the following directory location so that Message Analyzer can locate and load it during startup: `%LocalAppData%\Microsoft\MessageAnalyzer\OpnAndConfiguration\TextLogConfiguration\DevicesAndLogs\` Thereafter, whenever you use the **Add Files** feature in a Data Retrieval Session to load data from a target text log into Message Analyzer, the drop-down list in the **Text Log Configuration** column on the **Files** tab of the **New Session** dialog will contain your new configuration file as a list item. This enables you to select this configuration file whenever it is required to parse the unique text log messages for which it is designed.

Specifying a Default Configuration File for All Text Logs

You have the option to specify a particular OPN configuration file that will be used by default for all text log data files from which you load data into Message Analyzer. You can do this from the global **Options** dialog that is accessible from the Message Analyzer **Files** tab. On the **General** tab of the **Options** dialog, click the **Default text log configuration** drop-down list and select one of the predefined text log configuration files as your default, or you can select a custom configuration file that you created as your default. However, you will need to have placed your custom configuration file in the previously specified location in `%LocalAppData%` before it will appear in the drop-down list.

Your selection in the **Text Log Files** pane of the **Options** dialog for the default text log configuration file updates the `<DefaultLogFileReaderModule>` section of the app.ApplicationConfiguration.cfg file in the following directory

location: `%LocalAppData%\Microsoft\MessageAnalyzer\` Thereafter, whenever you use the **Add Files** feature on the **Files** tab of the **New Session** dialog to target a text log from which to load data in a Data Retrieval Session, the default OPN configuration file that you selected in the **Options** dialog is automatically selected in the drop-down list of the **Text Log Configuration** column, at the time when the actual text log data file appears in the files list. Moreover, if you launch a text log data file from **Windows Explorer**, **Quick Open**, or you use the drag-and-drop method when a default OPN configuration file has been specified, Message Analyzer will automatically begin loading data, at which point your text log messages begin to accumulate in the default **Analysis Grid** viewer.

> **NOTE**
>
> If you want to override the default configuration file specification, you can simply select the **(None)** item from the drop-down list in the **Text Log Files** pane of the **Options** dialog.

Addendum 2: HTTP Status Codes

3 minutes to read

This section contains reference information for the HTTP protocol that you can use to troubleshoot HTTP messages with Message Analyzer.

HTTP Status Codes

The table that follows describes common HTTP client and server **StatusCodes** and associated **ReasonPhrases** that can be returned from HTTP message exchanges between client browsers and web servers. Definitions for these codes and phrases are widely known to most networking audiences; however, they are reproduced here for user convenience, to provide the benefit of having this information readily available when performing various procedures in this documentation.

Table A1. HTTP StatusCode and ReasonPhrase Definitions

STATUSCODE	REASONPHRASE	DEFINITION
Common Client Errors		
400	Bad Request	Indicates a syntax error in the request, which is therefore denied. This is basically an error message from the web server indicating that the web browser incorrectly attempted to access a resource or the request was corrupted.
401	Authorization Required	Indicates that the request header did not contain the necessary authentication codes and therefore the requesting client is denied access. This error usually occurs when a website visitor attempts to access a restricted webpage but does not have authorization to do so, often because of a login failure.
403	Forbidden	Can indicate that a client is not permitted to view a particular file or has attempted to access a forbidden directory. This StatusCode is also returned when the web server is unable to accommodate more visitors. Note that this error is similar to the 401 authentication error, although in the case of the 403 error, there was no login opportunity available.
404	Not Found	Indicates that a requested resource was not found on the server, typically because it does not exist. This error can also be caused by the invalid spelling of a URL, a broken link, or the resource was moved to another location.

STATUSCODE	REASONPHRASE	DEFINITION
405	Method Not Allowed	Indicates that the method being used by a requestor to access a resource is not permitted.
406	Not Acceptable	Indicates that the requested resource exists but cannot be returned to the client system because of an incompatible format.
407	Proxy Authentication Required	Indicates that the request must be authorized before proceeding any further.
408	Request Timed Out	Often the result of heavy network traffic, which could indicate that a web server is overburdened.
409	Conflicting Request	Might be that there are too many concurrent requests for a resource, which can be indicative of the inability of a server to handle the amount of requests it is receiving.
410	Gone	Indicates that a resource previously existed in a particular location, but is no longer present there.
411	Content Length Required	Indicates that the request is missing the required Content-Length header.
412	Precondition Failed	Indicates that the client does not have the required configuration set up for a resource to be delivered by the web server.
413	Request Entity Too Long	Indicates that the requested resource is too large to process.
414	Request URI Too Long	Indicates that the entered URI address is too long for the server.
415	Unsupported Media Type	Indicates that the file type specified in the request is unsupported.
Common Server Errors		
500	Internal Server Error	Could indicate an overloaded web server that cannot properly handle the requests it is receiving.
501	Not Implemented	Indicates that the request cannot be performed by the web server.

STATUSCODE	REASONPHRASE	DEFINITION
502	Bad Gateway	This error usually indicates one or more of the following: - Improperly configured proxy servers. - Faulty IP communication between back-end computer nodes. - The client's ISP is overloaded. - A firewall is improperly functioning. Typically can be resolved by clearing the client cache to force a different proxy server into use to resolve the web server content.
503	Service Unavailable	Indicates that the requested service, file, or resource is currently unavailable.
504	Gateway Timeout	Indicates that the server gateway has timed out, which could indicate that one or more servers are over-burdened by heavy network traffic. Time-outs can occur when a server in a chain of servers does not receive a timely response from another server, indicating slow communication between upstream computers.
505	HTTP Version Not Supported	Indicates that the server does not support the HTTP version in use for the client request.