RSA® NetWitness

Version 11.7

NetWitness Investigate User Guide

Contact Information

RSA Link at https://community.rsa.com contains a knowledge base that answers common questions and provides solutions to known problems, product documentation, community discussions, and case management.

Trademarks

RSA Conference Logo, RSA, and other trademarks, are trademarks of RSA Security LLC or its affiliates ("RSA"). For a list of RSA trademarks, go to https://www.rsa.com/en-us/company/rsa-trademarks. Other trademarks are trademarks of their respective owners.

License Agreement

This software and the associated documentation are proprietary and confidential to RSA Security LLC or its affiliates are furnished under license, and may be used and copied only in accordance with the terms of such license and with the inclusion of the copyright notice below. This software and the documentation, and any copies thereof, may not be provided or otherwise made available to any other person.

No title to or ownership of the software or documentation or any intellectual property rights thereto is hereby transferred. Any unauthorized use or reproduction of this software and the documentation may be subject to civil and/or criminal liability.

This software is subject to change without notice and should not be construed as a commitment by RSA.

Third-Party Licenses

This product may include software developed by parties other than RSA. The text of the license agreements applicable to third-party software in this product may be viewed on the product documentation page on RSA Link. By using this product, a user of this product agrees to be fully bound by terms of the license agreements.

Note on Encryption Technologies

This product may contain encryption technology. Many countries prohibit or restrict the use, import, or export of encryption technologies, and current use, import, and export regulations should be followed when using, importing or exporting this product.

Distribution

Use, copying, and distribution of any RSA Security LLC or its affiliates ("RSA") software described in this publication requires an applicable software license.

RSA believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." RSA MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

© 2020 RSA Security LLC or its affiliates. All Rights Reserved.

January 2022

Contents

How NetWitness Investigate Works 14

Metadata, Meta Keys, Meta Values, and Meta Entities 14

Triggers for an Investigation 15

Workflow of an Investigation 15

Focus on Metadata, Query, and Time 23

Focus on Respond View Incidents and Alerts 24

NetWitness Investigate Views 24

Navigate View 24

Events View 25

Legacy Events View 27

Contextual Information for an Event 28

Reconstruction and Event Analysis 30

Configuring NetWitness Investigate Views and Preferences 32

Configure the Navigate View and Legacy Events View 33

Configure Common Settings 33

Access the Navigate View and Legacy Events View Settings 34

Calibrate Navigate View Value Loading Parameters 36

Configure Navigate View and Legacy Events View Parameters 37

Configure the Default Log Export Format 38

Configure the Default Meta Value Export Format 38

Calibrate Legacy Events View Retrieval and Default Reconstruction 38

Enable or Disable Cascading Style Sheet Rendering in Web Content Reconstructions 39

Configure Search Options 39

Configure the Events View 41

Set the Default Investigate View 41

Set User Preferences for the Events View 42

Beginning an Investigation 46

Focus on Metadata, Raw Events, and Event Analysis 46

Focus on Hosts and Files 47

Focus on Risky User and Entity Behavior 47

Focus on Scanning Files for Malware 47

Begin an Investigation in the Navigate or Legacy Events View 48

Begin an Investigation (No Default Service) 48

Set or Clear the Default Service 50

Begin an Investigation (Default Service Specified) 51

Change the Service or Collection to Investigate 52
Investigate Workbench Restoration Collections 53
Begin an Investigation in the Events View 55
Access the Events View 56
Access the Events View (Version 11.0) 58
Refining the Results Set 59
Use Meta Groups to Focus on Relevant Meta Keys 60
Live Meta Groups 60
Default Meta Keys Group (Version 11.5 Events View) 62
Custom Meta Groups 62
Work with Meta Groups in the Events View (Version 11.5 and Later) 63
View the Meta Keys in a Meta Group 63
Select a Meta Group 66
Create a Custom Meta Group 67
Delete a Custom Meta Group 70
Edit a Custom Meta Group 71
Copy a Meta Group (Version 11.5 and Later) 74
Meta Group Folders 77
Create a Meta Group Folder 78
Edit and Move Meta Group Folder 79
Copy Meta Group Folder 79
Copying Meta Group Private or Shared Folders 80
Copy Meta Group Folder Deployed from Live 81
Delete Meta Group Folder 83
Work with Meta Groups in the Navigate View 84
Create a Meta Group and Add Meta Keys 84
Copy and Edit a Meta Group 87
Edit a Custom Meta Group 87
Delete a Meta Group 89
Export a Meta Group 89
Import a Meta Group 89
Use Columns and Column Groups in the Events List 91
Built-In Column Groups 96
Live Column Groups 96
Custom Column Groups 100
Filtering Folders 101
Dialogs for Managing Column Groups 102
Work with Columns and Column Groups in the 11.4 and Later Events View 103
Manually Select Columns to Display and Adjust Column Order and Width 104
Select a Column for Sorting Events in the Events Panel (Version 11.4) 105

Sorting by Column (Version 11.4.1 and Later) 105
Sorting by Column (Version 11.4) 107
View the Meta Keys Included in a Column Group 108
Select a Column Group 110
Create a Custom Column Group 112
Delete a Custom Column Group 116
Edit a Custom Column Group 118
Create a Copy of a Column Group (Version 11.5 and Later) 121
Create a Column Group Folder 124
Edit and Move Column Group Folder 125
Copy Column Group Folder 125
Copy Group Folder Deployed from Live 126
Delete Column Group Folder 127
Select a Column Group and Columns (11.3 and Earlier Event Analysis View) 128
Work with Column Groups in the Legacy Events View 129
Select a Column Group 129
Create a Custom Column Group in the Legacy Events View 130
Delete a Column Group (Legacy Events View) 133
Edit a Column Group (Events View) 134
Import and Export a Column Group (Legacy Events View) 137
Use Query Profiles to Encapsulate Common Areas for Investigation 139
Built-In Query Profiles 139
Live Query Profiles 140
Custom Query Profiles 140
Dialogs for Managing Query Profiles 141
View Query Profile Details (Events View) 143
Apply a Query Profile (Events View) 145
Create or Edit a Custom Query Profile (Events View) 146
Delete a Custom Query Profile (Events View) 150
Copy a Query Profile (Version 11.5 and Later) 152
Create a Query Profile Folder 154
Edit and Move Query Profile Folder 154
Copy Query Profile Folder 155
Copy Query Profiles Group Folder Deployed from Live 156
Delete Query Profile Folder 156
Navigate to the Manage Profiles Dialog (Navigate and Legacy Events Views) 157
Create, Edit, or Delete a Profile Group (Navigate or Legacy Events View) 158
Create and Edit Profiles (Navigate or Legacy Events View) 160
Delete a Profile (Navigate or Legacy Events View) 161
Change the Active Profile (Navigate or Legacy Events View) 161

Import Profiles (Navigate or Legacy Events View) 162
Download Profiles (Navigate or Legacy Events View) 162
Drill into Metadata in the Events View (BETA) 163
Modes of Operation 164
View Metadata in the Filter Events Panel 165
Show Max Value of Meta Groups 167
View the Context Lookup Panel in the Filter Events Panel 168
Understand Visible Metadata 169
Stop and Resume Metadata Loading 170
Close All Except One Meta Key (Version 11.5.1 and Later) 171
Set the Ordering Method for Meta Values 172
Drill into Meta Values 174
Copy the Meta Values for a Meta Key 176
View a Selected Meta Value in RSA Live 178
Refocus the Investigation of a Meta Value 178
Filter Results in the Events View 181
Initial Filter Using the Query Bar 181
Find a Text String in the Events Panel (Version 11.4 and Later) 182
Refining the Results in the Events Panel 183
Query Builder Concepts 184
Guided Mode vs. Free-Form Mode 185
Concepts for Editing Multiple Filters 186
The Version 11.3 and Earlier Query Builder 188
The Version 11.4 Query Builder 188
Meta Keys Cached for Faster Loading 188
Text Filter 189
Pasting Text Instead of Typing 189
Select All Filters and Copy All Filters (Version 11.4.1) 189
Use of Recent Queries 189
Use of Advanced Operators 189
Easy Use of AND/OR Operators 190
Automatically Balanced Parentheses 190
Hints about Available Values 190
CIDR Notation and Shorthand 191
Ranges or Series of Values 191
No Separating Space Required After Meta Key and Operator (Version 11.4.1) 191
Select a Time Range 191
Submit a Query 193
Cancel Execution of a Query 193
View Status of a Query 194

Build a Query in Guided Mode 196
Keyboard Actions to Use in Guided Mode 196
Visual Feedback in Guided Mode 199
Add a Simple Filter in Guided Mode 201
Add a Free-Form Filter in Guided Mode (Version 11.3 and Later) 205
Add a Text Filter to Find a Value Anywhere in the Data Set (Version 11.4 and Later) 206
Select All Filters and Copy All Filters in the Query Bar (Version 11.4.1 and Later) 209
Paste Text in the Query Bar (Version 11.4 and Later) 209
Insert a Filter Based on a Recent Query (Version 11.4 and Later) 211
Edit a Filter in Guided Mode 212
Query Using Selected Filters in Guided Mode 213
Delete a Filter and Delete Text or Parentheses in a Filter in Guided Mode 214
Create a Query in the Free-Form Mode 215
Filter Results in the Navigate View 217
Set the Time Range 217
Set the Quantification Method and Sort Sequence of Meta Key Results 219
Manage and Apply Default Meta Keys in an Investigation 220
Drill into Data in the Navigate View Time Chart 222
Drill into Data in the Values Panel 223
Filter Results in the Legacy Events View 231
Filter Events Displayed in the Legacy Events View 231
Page Through Events in the Legacy Events View 232
Create a Query in the Navigate and Legacy Events Views 233
Create a Query Using the Basic Method 233
Create a Query Using the Advanced Method 234
Apply a Recent Query 236
Search for Text Patterns in the Navigate and Legacy Events Views 238
Keyword Text Search 238
Options Controlling Search Behavior 239
Regular Expression Search Syntax 241
Raw Text Keyword Search 241
Search Procedures 241
Search in the Navigate View 241
Search in the Legacy Events View 241
View and Modify Queries Using URL Integration 242
Service Id Known 242
Host and Port Known 242
Examples 243
Additional Notes 243

Reconstructing and Analyzing Events **244**
Examine Event Details in the Events View 247
Event Details for Each Event Type 247
Text Reconstruction 248
Packet Reconstruction 252
File Reconstruction 254
Host Information 255
Email Reconstruction 258
Analyze Events in the Events View 261
How Results Are Loaded and Sorted 261
Actions to Refine the Events List 262
Actions to Analyze Events 263
Open, Close, and Adjust the Size of the Panels in the Events View 263
Select the Analysis Type for an Event 264
Adjust the Display of Requests and Responses 264
View Associated Metadata for an Event 265
Show or Hide the Event Header 268
Page Through Events in the Packet and Text Tabs 269
Expand Truncated Text Entries in the Text Tab 270
Perform URL and Base64 Encoding and Decoding in the Text Tab 270
View Decompressed Text in an HTTP Network Session in the Text Tab 272
View a JSON String in Tree Format in the Text Tab (Version 11.5.1) 274
Use the Payload Only Option in the Packet Tab 275
View Highlighted Bytes in the Packet Tab 276
Highlight Common File Types in the Packet Tab 277
Reconstruct an Event in the Legacy Events View 280
Reconstruct an Event Using an Event ID 280
Reconstruct an Event from a Drill Point in the Navigate View 281
View Side by Side or Top to Bottom 283
Select Event Information to View 283
Select Event Reconstruction Type 283
Open or Download an Email Attachment 284
Export an Event as a PCAP File 284
Extract Files from a Reconstructed Event 285
Look Up Additional Context for Results 286
Open the Context Lookup Panel 287
Add an Entity to a Whitelist 290
Create a List (Events View) 290
Pivot to Investigate > Navigate (Events View) 291
Pivot to Archer (Events View) 291

Pivot to NetWitness Endpoint Thick Client (Events View) 292
View the Context Lookup Panel in the Navigate View or Legacy Events View 292
Add Meta Values to an Existing List (Navigate and Legacy Events Views) 293
Remove a Meta Value from a Context Hub List (Navigate and Legacy Events Views) 294
Create a New List (Navigate and Legacy Events Views) 294
Launch a Lookup of a Meta Key 296
Launch an Endpoint Thick Client Lookup in the Events View 296
Launch an Endpoint Thick Client Lookup in the Navigate View 297
Perform Lookups of Meta Values in Events 299
Launch Other External Lookups from the Navigate View 301
Launch a Malware Analysis Scan from the Navigate View 303
Group Events from Split and Related Sessions in the Events and Legacy Events Views 305
Split Network Session 306
Session Size and Time Split 306
Transaction Handling Split 307
Session Fragments Highlighting 308
Related Network Session 308
Use Cases for Viewing Events from Split and Related Network Sessions 309
Show and Hide Relationships in the Events List 310
Session Fragments Highlighting (11.3 Events View) 311
Find and Combine Fragments in the Legacy Events View 312
Visualize Metadata as Parallel Coordinates 315
Best Practices for Effective Parallel Coordinates Charts 315
RSA Meta Groups for Parallel Coordinates Use Cases 316
View a Parallel Coordinates Visualization 316
Select Meta Keys for a Parallel Coordinates Visualization 317
Optimize a Parallel Coordinates Visualization 322
Sample Use Case 324
Sample Visualization of a Large Data Set 325
Visualize the Current Drill Point in Informer 327
Downloading and Acting Upon Results 328
Download Data in the Events View 329
Download Events or Metadata in the Events Panel 329
Download a Log in the Text Reconstruction 332
Download Network Event Data in the Text or Packet Reconstruction 333
Download Files from a Network Event in the File Reconstruction 335
Download Attachments from an Email Reconstruction 336
Export or Print a Drill Point in the Navigate View 340
Export Events in the Legacy Events View 342
Add Events to an Incident in the Events View 343

Add Events to an Incident in the Legacy Events View 345
Troubleshooting NetWitness Investigate 347
PCAPS and SDK Issues 347
Navigate View and Legacy Events View Issues 349
Events View Issues 350
Investigate Reference Materials 356
Add Events to an Incident Dialog 357
Workflow 357
What do you want to do? 357
Related Topics 358
Quick Look 359
Add/Remove from List Dialog 361
Workflow 361
What do you want to do? 362
Related Topics 362
Quick Look in the Events View 363
Quick Look in the Navigate and Legacy Events Views 364
Column Groups Dialogs 366
Related Topics 367
Quick Look - Column Group Menu, Create Column Group Dialog, and Column Group Details Dialog 367
Quick Look - Manage Column Groups Dialog 371
Context Lookup Panel 373
Workflow 373
What do you want to do? 373
Related Topics 374
Quick Look (in the Navigate and Legacy Events Views) 375
Incidents 376
Alerts 376
Lists 377
Endpoint 377
Quick Look in the Events View (Version 11.2 and Later) 377
Lists Tab 380
Archer Tab 381
Active Directory Tab 382
NetWitness Endpoint Tab 383
Alerts Tab 385
Incidents Tab 386
Live Connect Tab 388
File Reputation Tab 393

TI Tab 394
REST API Tab 395
Create an Incident Dialog 397
Workflow 397
What do you want to do? 397
Related Topics 398
Quick Look 398
Events View 400
Workflow 400
What do you want to do? 400
Related Topics 401
Quick Look 402
Filter Events Panel 405
Query Console 406
Timeline 407
Quick Look for Version 11.0.0.x (End-Of-Life) 408
Events View - Email Tab 410
Workflow 410
Related Topics 410
Quick Look 410
Events View - File Tab 412
Workflow 412
What do you want to do? 412
Related Topics 413
Quick Look 413
Events View - Host Tab 415
Workflow 415
What do you want to do? 415
Related Topics 416
Quick Look 416
Events View - Packet Tab 418
Workflow 418
What do you want to do? 418
Related Topics 419
Quick Look 419
Events View - Text Tab 421
Workflow 421
What do you want to do? 421
Related Topics 422
Quick Look 422

Investigate Dialog 424
Workflow 424
What do you want to do? 424
Related Topics 425
Quick Look 426
Investigation Tab - User Preferences Panel 428
Related Topics 428
Quick Look 428
Investigate View 432
Legacy Event Reconstruction View 433
What do you want to do? 433
Related Topics 434
Quick Look 434
Legacy Events View 437
What do you want to do? 437
Related Topics 438
Quick Look 438
Detailed Description 441
Manage Default Meta Keys Dialog 444
Related Topics 444
Quick Look 444
Meta Groups Dialogs 446
Related Topics 446
Quick Look - Meta Groups Menu, Create Meta Group Dialog, and Meta Group Details Dialog 446
Quick Look - Manage Meta Groups Dialog 449
Navigate View 452
Workflow 452
What do you want to do? 453
Related Topics 453
Quick Look 454
Toolbar 454
Pause/Reload Button and Breadcrumb 457
(Optional) Debug Information 458
Time Banner 458
Visualizations 458
Timeline Chart 458
Parallel Coordinates Chart 460
Values Panel 462
Values Panel Loading Behavior 463
Iterative results 464

Partial results 464
Debug Information 465
Load Complete 465
Query Dialog 466
What do you want to do? 466
Related Topics 467
Quick Look 467
Simple View 467
Advanced View 468
Recent View 468
Query Profiles Dialogs 470
Related Topics 470
Quick Look - Query Profile Menu, Create Query Profile Dialog, and Query Profile Details Dialog 471
Quick Look - Manage Profiles Dialog 473
Settings Dialogs for Investigate Views 475
Related Topics 475
Quick Look 475
Navigate View Settings Dialog 476
Legacy Events View Settings Dialog 477
Events View Preferences Dialog 479

How NetWitness Investigate Works

NetWitness Investigate provides analysts the means to analyze events that have been captured by NetWitness. Using Investigate, analysts can examine packet, log, and endpoint data, and identify possible internal or external threats in their environment. There are several different views available to analysts to gain different perspectives into the data in their environment. A key element that all the views have in common is metadata.

Note: In Version 11.1 and later, the Hosts and Files views provide a view into endpoint data. Earlier versions offer access to endpoint data using a standalone NetWitness Endpoint server.

Metadata, Meta Keys, Meta Values, and Meta Entities

RSA NetWitness audits and monitors all data communications in an environment. One type of service--a Decoder--ingests, parses, and stores the original packets captured on the network, logs forwarded by a device, and endpoint events seen by the endpoint agent. The configured rules, parsers, and feeds on the Decoder create *metadata* that analysts can use to investigate the ingested logs, packets, and endpoint data. Another type of service, called a Concentrator, indexes and stores the metadata, making it more efficient to search through all types of metadata.

The metadata is created to give analysts valuable points of reference associated with the original data. This allows analysts to quickly get a sense of what has transpired without being required to examine every detail of an event. The metadata is in the form of a *meta key* and *meta values* for the key. For example, `ip.src` is a meta key, and an IP address (192.168.1.1) that is the source of the traffic is a meta value tagged as `ip.src`. When you view data in Investigate, you see the meta key `ip.src` and all of the IP addresses (meta values) that are tagged with that key. Some meta keys are built-in and others may be custom keys specific to your environment and defined by the administrator. All metadata, no matter the source of the data, is normalized into the Unified Data Model for RSA NetWitness Platform to keep similar metadata concepts grouped together into like meta keys (see https://community.rsa.com/community/products/netwitness/rsa-content/udm).

Meta entities are available in Version 11.1 and later. A *meta entity* is an alias that groups together the results from other meta keys. Meta entities organize similar meta keys into a single, easier to use, meta type. For example, the default Core database language includes distinct meta keys for IP source and IP destination. One of the built-in meta entities named `ip.all` represents the combined set of all IP sources and destinations. Some meta entities are already included by default, and the administrator can create custom meta entities. Analysts can use a meta entity in a query, a meta group, a column group, and a query profile. Parallel coordinates visualizations do not support meta entities. Administrators can use meta entities to define a query prefix to apply to a user role and a user as described in the *System Security and User Management Guide*. The *Decoder Configuration Guide* provides additional information about creating meta entities and how they can be used in rules.

Note: Meta entities need to be configured on all upstream Concentrators. If any Concentrator does not have a meta entity configured, that meta entity will be empty when you query the Broker.

Analysts usually query the Broker or Concentrator to discover threats. The Concentrator handles queries, only going to the Decoder for raw logs or endpoint events or a full reconstruction of network events. ESA, Malware Analysis, and Reporting Engine also query the Concentrator, where they can quickly get all the pertinent metadata associated with an event and generate information about the event without having to query each Decoder. In some special cases, analysts may query a Decoder.

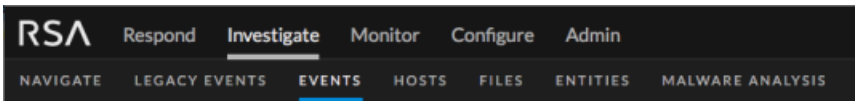
Triggers for an Investigation

These are a few examples of triggers for an investigation:

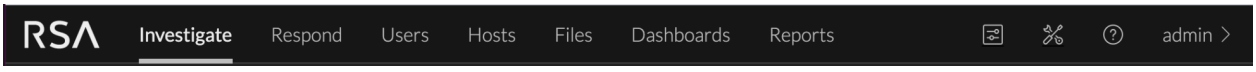
- You receive intelligence about a new active directory hack. Starting in the Events view, you use that intelligence to run a search across all of your raw Active Directory log data for the last 24 hours.
- You are asked by the SOC manager to find any Pokemon Go malware due to its popularity. Starting in the Navigate view, you craft a query to look for an HTTP session using a specific user agent related to the malware that your SOC manager found on a security blog.
- An incident responder escalates a ticket that shows some odd indicators related to a host. Starting in the Hosts view, you examine that host to find specific details.
- You are looking for the next zero day attack and start drilling into the network metadata in the Navigate view (or the Filter Events panel in the Events view) to find any abnormal automated sessions leaving the enterprise.
- You are asked by your SOC manager to find any information related to user `jarvis`, an employee who was just let go. Starting in the Users view you can filter for that username to make sure there is no longer any activity for that user and see if that user deviated from their typical behavior prior to being let go.
- A phishing attack detected has an associated attachment, and you want to know what devices in your environment have seen that file by searching for the file hash in the Files view.
- A malicious file has been automatically found in your environment, and you want to review the static and dynamic analysis done on that file along with how many systems it has been transmitted to or from. Starting in Investigate > Malware Analysis you can see the analysis results.

Workflow of an Investigation

Analysts can investigate data captured by NetWitness, and deep dive from information on a NetWitness dashboard, the Springboard (Version 11.5 and later), a NetWitness Respond incident or alert, a report created by the NetWitness Reporting Engine, or a third-party application. During the course of an investigation, analysts can move seamlessly between views: the Navigate view, the Events view, the Legacy Events view, the Hosts view, the Files view, the Users (Entities) view, and the Malware Analysis view. By default, the Navigate view and the Legacy Events view are disabled. This figure illustrates the NetWitness Investigate submenus in Version 11.4 and earlier.



This figure illustrates the Version 11.5 menu, which is optimized for the analyst workflow with Users, Hosts, and Files moved to the top-level menu.



Note:
- The Files and Hosts views are available in Version 11.1 and later. Before Version 11.5, they were a submenu of Investigate.
- The Users view is available in Version 11.2 and later; in Version 11.4 it is labeled Entities view. Before Version 11.5, it was submenu of Investigate.
- By default, the Legacy Events view is disabled in Version 11.4, but can be enabled by an administrator as described in the *System Configuration Guide*.
- By default, the Navigate view is disabled in Version 11.6 as the Filter Events Panel in the Events view provides this functionality. To enable the Navigate view, see Configure the Navigate View and Legacy Events View.
- Specific user roles and permissions are required for a user to conduct investigations and malware analysis in NetWitness. If you cannot see a view, the administrator may need to adjust the roles and permissions configured for you.

The views are closely integrated to reduce the need to jump from one view to another. Each use case determines the starting point for your investigation; every situation is unique in terms of the types of information you are attempting to find. Many investigations start in one view, and end in a different view as you learn something and then need to follow that result to a different line of questioning. Experienced analysts frequently begin an investigation in the Navigate view or the Events view. Less experienced analysts can begin in the Dashboard, Respond view, or the Springboard (Version 11.5 and later), in which clickable incidents and alerts link to detailed information and analyses in the different views.

Go to...	Focus
Navigate view	The Navigate view (Version 11.5 and earlier) works well for providing a high-level vantage point of what has been seen in your environment for a specific time range by showing meta keys and meta values for log, endpoint, and network events. After drilling into meta values, go to the Events view to see the raw event (See Filter Results in the Navigate View.)

Go to...	Focus
Events view	The Events view (default Investigate view) is the workflow for analysts interacting with events, presenting different facets of the same data in adjacent panels. In Version 11.6, there is no need to go to the Navigate view to drill into meta values as the Filter Events Panel in the Events view provides this functionality. (See Refining the Results Set, Reconstructing and Analyzing Events, and Downloading and Acting Upon Results.)

Go to...	Focus
Legacy Events view	The Legacy Events view was the original workflow for looking at event details in Version 11.0 to 11.3.x.x. The Legacy Events is replaced by the 11.4 and later Events view and it is hidden unless the administrator enables it. (See Refining the Results Set, Reconstructing and Analyzing Events, and Downloading and Acting Upon Results.)

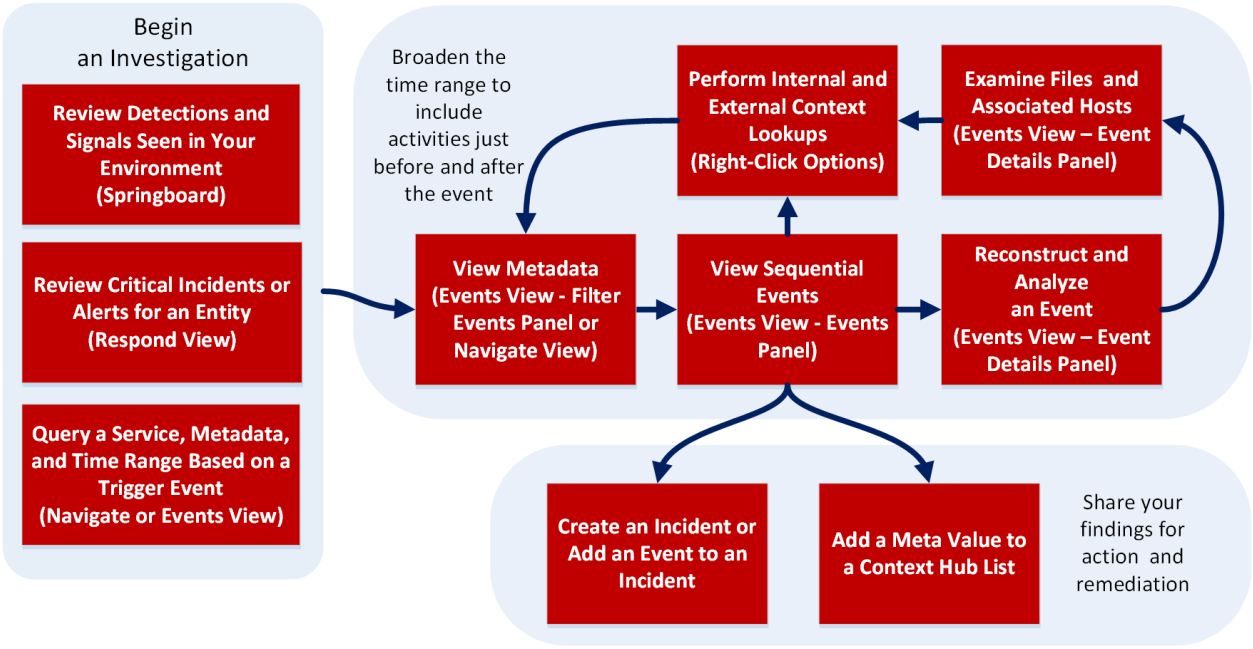
Go to...	Focus
Hosts view	The Hosts view option was moved to the main menu in Version 11.5. Hosts on which the NetWitness Endpoint agents are running are listed. For every host, you can view processes, drivers, DLLs, files (executables), services, anomalies, and autoruns that are running, and information related to logged-in users. (See the *NetWitness Endpoint User Guide*.)

Go to...	Focus
Files view	The Files view option was moved to the main menu in Version 11.5. Unique files in your deployment, such as PE, Macho, and ELF, are listed. For each file, you can view details such as file name, reputation status, file status, risk score, signature, checksum, and others. (See the *NetWitness Endpoint User Guide*.)
Malware Analysis view	If you are running a Malware Analysis appliance, you can scan files and see results of four types of analysis: network, static, community, and sandbox. If a file is malware, you can go to the Hosts view to see which hosts downloaded the file. (See the *Malware Analysis User Guide*.)

Go to...	Focus
Users view	The Users view (labeled as the Entities view in Version 11.4) was moved to the main menu in Version 11.5. It provides visibility into risky user behaviors across your enterprise using NetWitness UEBA. You can view a list of high-risk users and a summary of the top alerts for risky behavior for your environment, and then select a user or an alert and view details about the risky behavior, and a timeline during which the behaviors occurred. NetWitness Platform users assigned the Administrators or UEBA Analysts role have access to this view. (See the *NetWitness UEBA User Guide*.)

Focus on Metadata, Query, and Time

The following figure depicts the workflow for an investigation with focus on metadata, a query, and time range.



Analysts use Investigate to hunt for events that drive the incident response workflow and to do strategic analysis after another tool has generated an event. Beginning in the Navigate view, Events view, or Legacy Events view:

1. Start by executing a query on a service for a specific time range, filter results to get a subset of events, reconstruct or analyze an event, and repeat the process to reconstruct or analyze another event. Built-in query profiles, meta groups, and column groups provide a good starting point. For example, you can choose the RSA Email Analysis query profile to see only metadata that is useful when investigating email risks.
2. When an event bears a closer look, view the context around the event, and decide whether to create an incident or add the event to an incident. If you decide not to add the event to an incident, you can run another query to gain further insight, which starts again at the beginning of the workflow.
3. If you notice suspicious activity or files on a specific host in the network, gather additional information about the host and files found on the host in the Hosts and Files view, or in a standalone NetWitness Endpoint server.
4. If you find a file or event that potentially contains malware, do a Malware Analysis scan of the file or open Malware Analysis and start a scan of the service on which the event was seen.

Here is one simple use case: If there is a concern regarding suspicious traffic with certain countries, the Destination Country meta key reveals all destinations and the frequency of the contact. Drilling into those values yields the specifics of the traffic, such as the IP address of the originator and the recipient. Checking other metadata can expose the nature of attachments exchanged between the two IP addresses. When suspect IP addresses are identified, looking at the addresses in the Navigate view or Events view with a broader time range can provide clues about what happened before and after the event being investigated.

Another use case is to investigate an alert to discover a malicious insider in the network who is exfiltrating intellectual property or other sensitive data from a specific IP address. The investigation begins with this meta value: `Upload without change request followed by download alert`. Start in the Navigate view or Events view by filtering the values to the IP address during the time range in which the alert was generated. Alerts metadata shows risk indicators as meta values, and you can click on different meta values to filter the events list and then reconstruct the event. Next extract files and examine the files to understand what happened. With this information, you can filter on the same IP address and broaden the time range to see activities before and after the event.

Focus on Respond View Incidents and Alerts

An analyst who is working on an incident or an alert in Respond can open the incident in Investigate to do a deeper analysis of the event or alert.

- The workflow to respond to an incident typically begins in the Respond view, where the analyst who is investigating an incident needs to gather intelligence about the incident in Investigate. Hover over an underlined entity in an incident or alert, such as an IP address, and select the action **Pivot to Investigate**. The Events view opens and is filtered for the selected entity. Defined meta keys are queried and the captured packets, logs, and endpoint events are displayed in the Events view.
- If you find events that are relevant to the incident, add the events to the incident in Respond. You can also create a new incident in Respond based on one or more events found in Investigate.
- From the Incident Details view Indicators panel in Respond, open the Events view to get a better understanding of an indicator event.

NetWitness Investigate Views

This section provides a brief description and example of the Navigate view, Events view, and Legacy Events view, as well as the context information, event details and reconstructions available in those views. Refer to the *Malware Analysis User Guide* for information about features and functions of the Malware Analysis view.

Navigate View

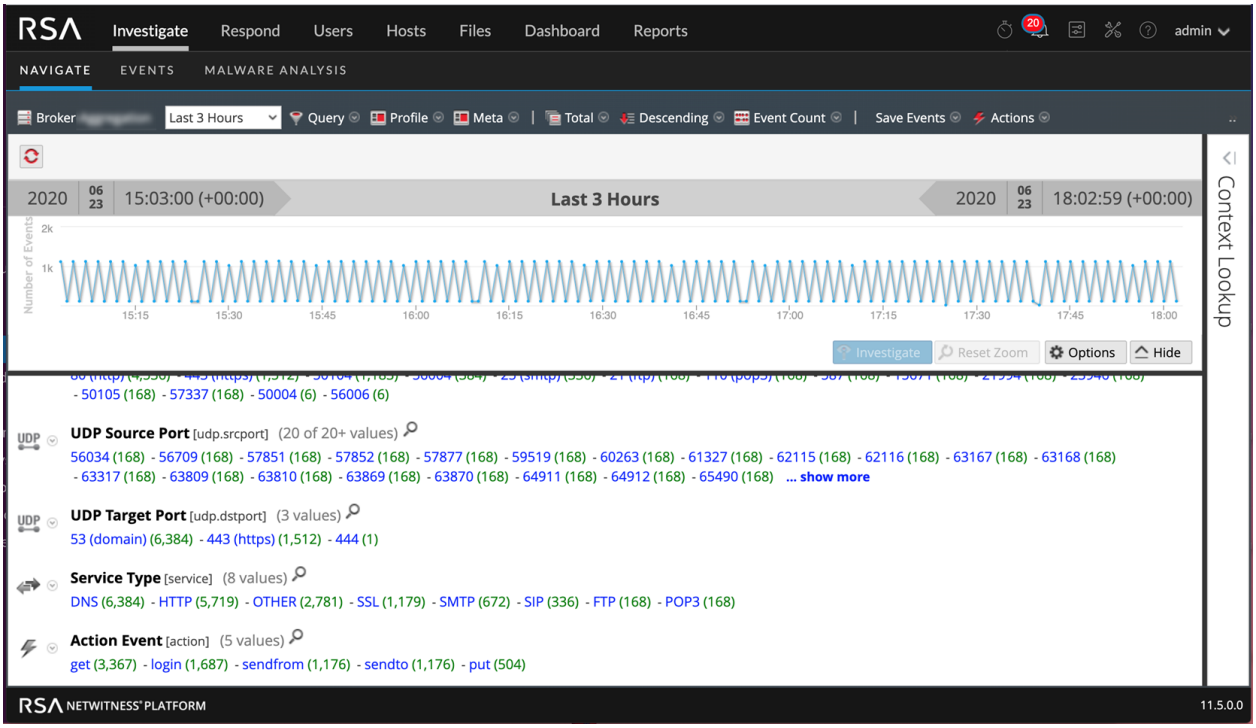
Note: By default, the Navigate view is disabled in Version 11.6 as the Filter Events Panel in the Events view provides this functionality. To enable the Navigate view, see Configure the Navigate View and Legacy Events View.

The Navigate view provides the capability to drill into and query metadata for network, logs, and endpoint events on a Broker, Concentrator, or Decoder (though investigating a Decoder is not typical). For certain configured meta keys, such as IP address, or hostname, you can see additional context information around a value using the Context Hub. The Navigate view also provides a sequential visualization of the data in a timeline. This figure illustrates the Navigate view.

- Each meta key listed displays the top 20 values based on how many events have those values.
- Drilling into the meta values by successive left or right clicking of values applies each clicked value as an additional filter to your query. As you drill, the subset of metadata seen grows smaller based on

the filters you have applied. For example, if you filter to only show HTTP (service=80), all the remaining metadata presented will be what is contained inside those HTTP events.

- With a smaller set of refined results, you can go to the Events view to examine further event details, or perform other lookups inside or outside the platform to gain further insight.



Events View

The Events provides the ability to view events sequentially, analyze raw event data and metadata, and (in Version 11.5) drill into metatdata as you can in the Navigate view.

- **In the Events panel**, Network, endpoint, and log events are listed in order by time. You can view the raw event, filter, sort, search, look at details and reconstructions, and download events. Clicking an event opens the Event Details panel for the event. Different reconstructions are available in the Events view (packets, text, files, email, and web) with helpful cues to identify points of interest, such as interesting bytes, file types, and encoded data.

- **In the Event Meta panel**, you can view related metadata for an event that is open in the Event Details panel. Analysts reviewing the metadata can change the order of the metadata to better track down what they are looking for. The items in the list of metadata can be grouped by the sequence they were generated or alphabetically.

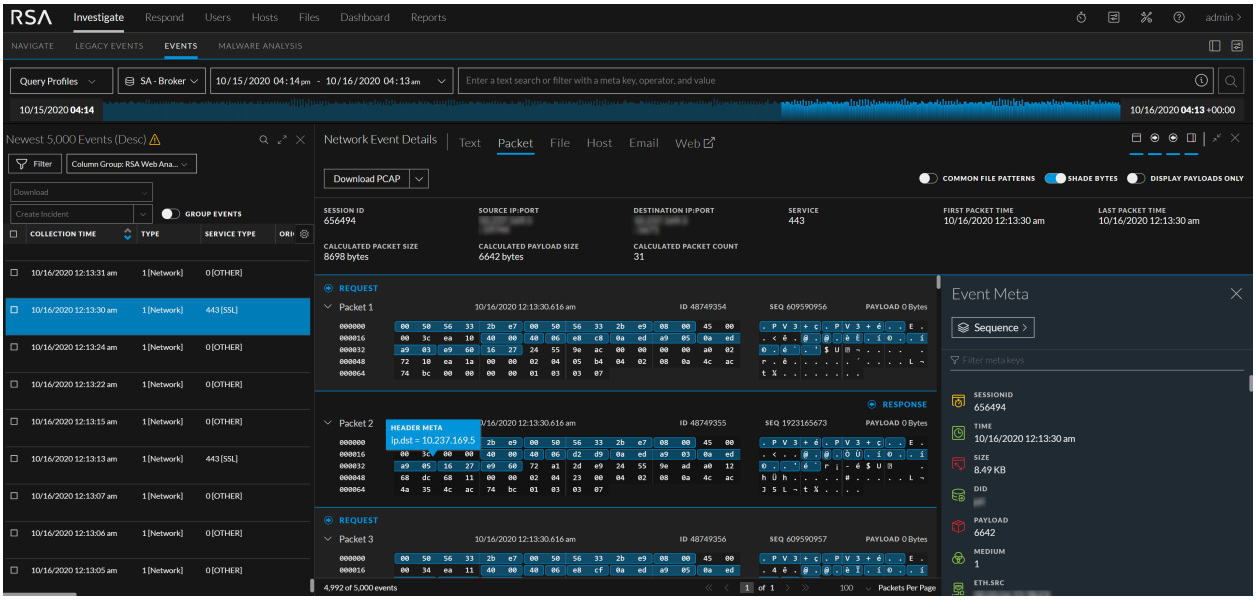
- **In the Filter Events panel** (Version 11.5 beta feature), you can drill into meta values for the listed events with your actions reflected in the Events panel. When the Filter Events panel is expanded to the full width of the browser, you can drill into meta values to hunt for specific information before

listing the events in the Events panel (comparable functionality to drilling into data in the Navigate view).
(Version 11.5.1) In the Filter Events panel, the meta values result threshold is 100000. If results are above the threshold, it is indicated using either ~ or >. For example, **(>100000)** indicates that the results are sorted based on count and are greater than the threshold. Similarly, **(~100000)** indicates that the results are sorted based on size and are greater than the threshold.

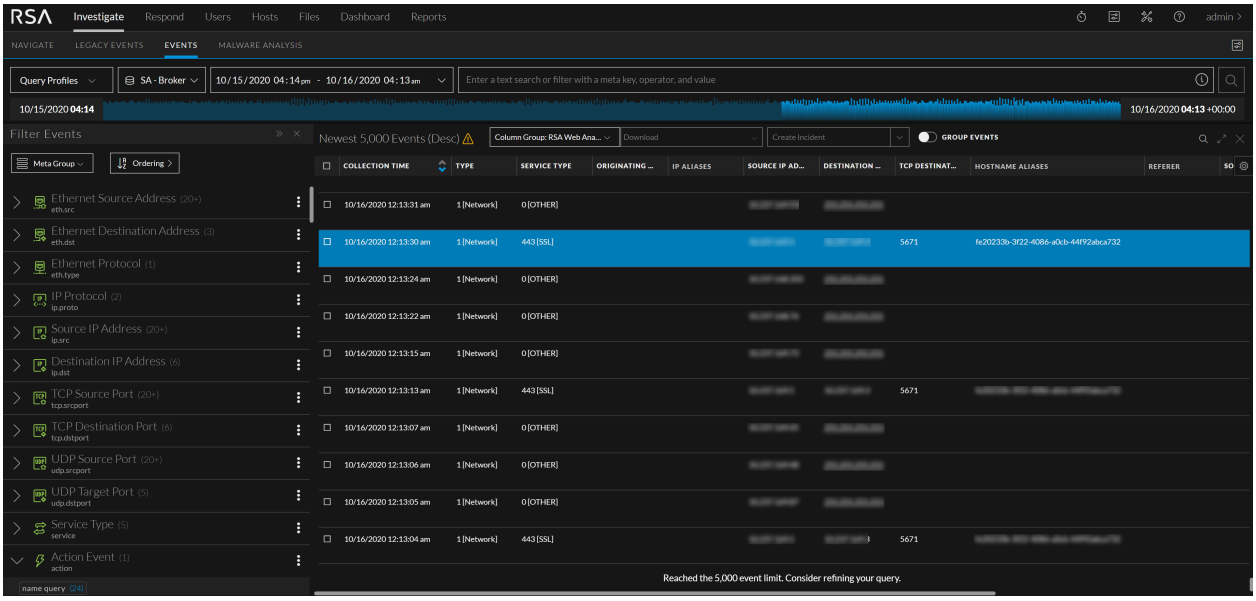
- **In the Event Details panel**, you can view details of a network, log, or endpoint event and safely reconstruct an event in a format similar to the original format. The tabs are Text, Packet, Files, Hosts, Users, Email, and Web.
- From various points in the Events view, you can pivot to standalone Endpoint, look up in Live, and do other internal lookups. External lookups allow you to search the internet for meta values with which you interacted, determine passive DNS information related to an IP address, check if a URL is blacklisted, and other third-party context integrations.
- (Version 11.5) If network data is enriched with endpoint data from an Endpoint deployment and the Endpoint Agent is configured for expanded network visibility, host information for network events is also displayed in the Events view in the header and in the Hosts tab.

  Note: For Expanded Network Visibility to work, ensure the service user account used for aggregating Endpoint Log Decoder data to Endpoint Concentrator is assigned with the `decoder.manage` permission. For more information on how to assign roles and permissions, see "Services Security View - Aggregation Role" in the *Hosts and Services Getting Started Guide for RSA NetWitness Platform*.
- For certain configured meta keys, such as IP address, or hostname, you can search for additional context information around a value using the Context Hub. The additional context may include incidents, alerts, Threat Intelligence, and other sources where the value was mentioned.
- You can export different types of data. When viewing files, you can export files in a zip archive to your local file system. When viewing email reconstructions, you can download attachments. You can download logs from the text reconstruction, and export packets from the packet reconstruction. You can download multiple events from the Events list.
- (Version 11.6) To view the download jobs, click  . Unlike the  icon in the **Legacy Events** and **Navigate** page, it does not open the Jobs tray. It opens the Jobs page where all jobs are listed. To display the Jobs tray, go to **Investigate** > **Navigate** (Version 11.5 and earlier) or **Investigate** > **Legacy Events** (Version 11.3 and earlier), and click the  (Jobs) icon.

This figure is an example of the Events view with a network event selected in the Events list, analyzed in the middle panel, and related metadata in the right panel.



This figure illustrates the Events view with the Filter Events panel open on the left side to filter the results set by drilling into meta values.

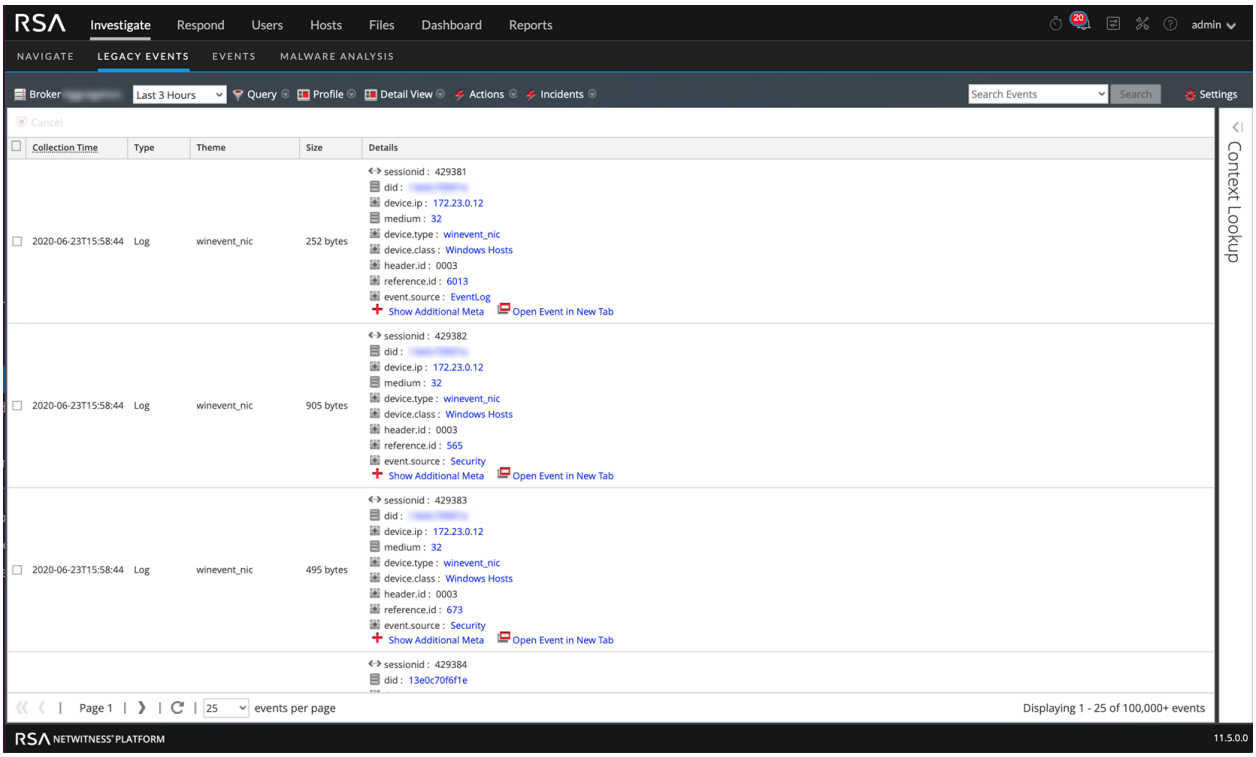


Legacy Events View

The Legacy Events view was the original user interface for analysts examining raw event data (11.0 to 11.3.x.x). The Legacy Events view is no longer needed in Version 11.4; it is hidden unless the administrator enables it as described under "Configure Investigation Settings" in the *System Configuration Guide*. When the Legacy Events view is enabled, both the Events view and the Legacy Events view are visible in the menu bar. The Legacy Events view provides a view of packet, log, and endpoint events in list form so that you can view events sequentially and reconstruct events safely.

- You can open the Legacy Events view for a meta value that you see in the Navigate view.
- For analysts without sufficient privilege to navigate a service, the Legacy Events view is a standalone investigation view in which analysts can access a list of network, log, and endpoint events from a NetWitness Core service without having to drill down through metadata first.
- The Legacy Events view presents event information in three standard forms, a simple list of events, a detailed listing of events, and a log view.
- For certain configured meta keys, such as IP address, or hostname, you can see additional context information around a value using the Context Hub. The additional context may include incidents, alerts, and other sources where the value was mentioned.
- You can export events and associated files, and create an incident from an event.

This figure illustrates the Legacy Events view.



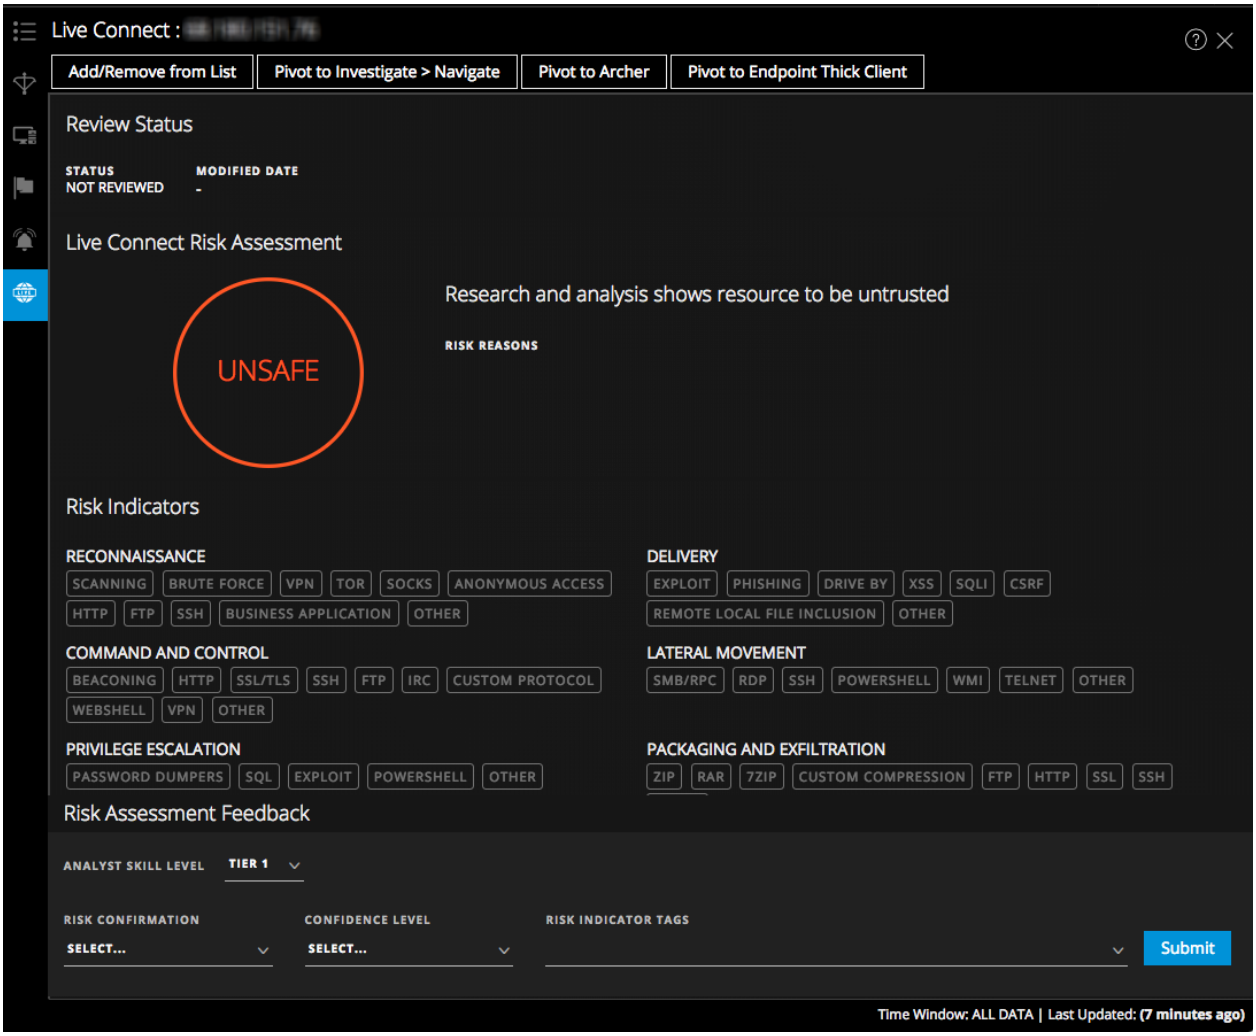
Contextual Information for an Event

In the Navigate view, Events view, and Legacy Events view, the Context Lookup panel shows details about elements associated with an event in the Context Hub for these meta types: IP Address, User, Host, Domain, MAC Address, Filename, and File hash. In addition, you can right-click all meta keys except time to see additional context.

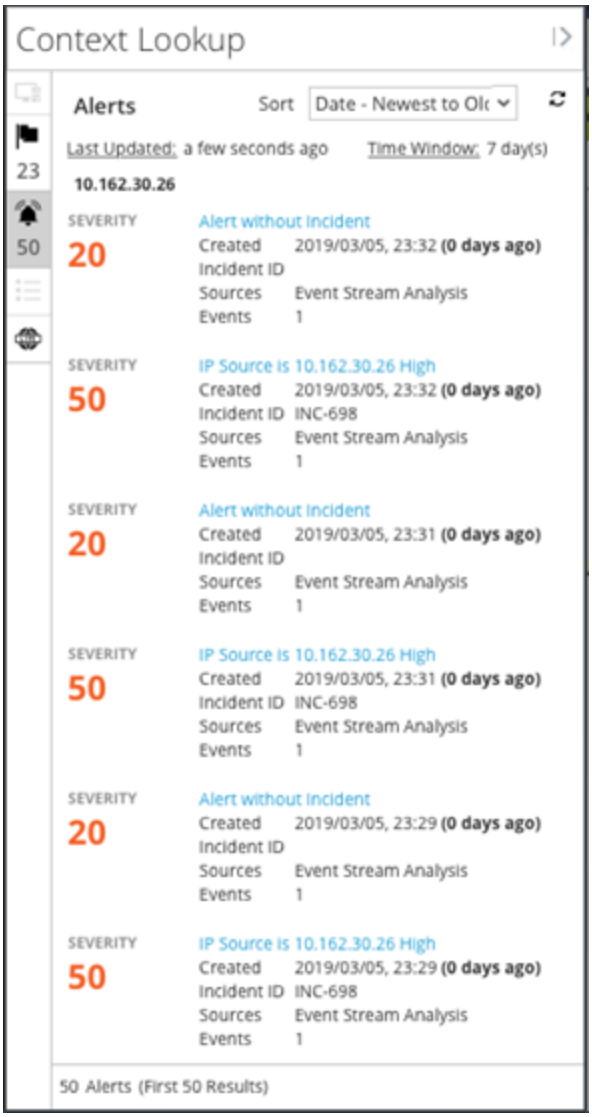
You can interact with the elements of an event to get further insight including related incidents, alerts, custom lists, Archer assets, active directory details, NetWitness Endpoint IIOCs, and STIX data sources (namely File, TAXII Server, and REST Server). (See Look Up Additional Context for Results.)

Note: Archer assets and active directory details are available in the Events view context lookup. Endpoint context lookup is available for NetWitness Endpoint 4.4.0.2 or later hosts, but not for the NetWitness Endpoint 11.1 or later hosts.

The following figure shows the Context Lookup panel, which opens to the right of the Events panel in the Events view.



The following figure shows the Context Lookup panel, which opens to the right of the Events list in the Legacy Events view.

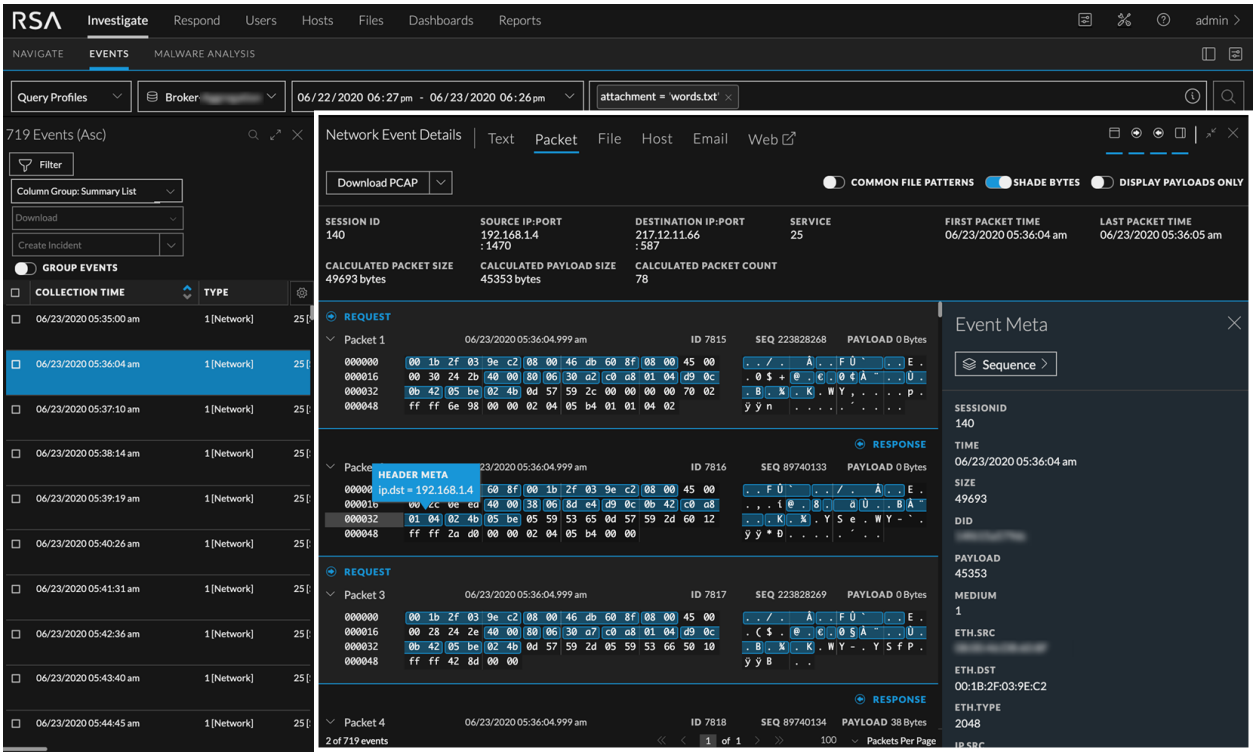


Reconstruction and Event Analysis

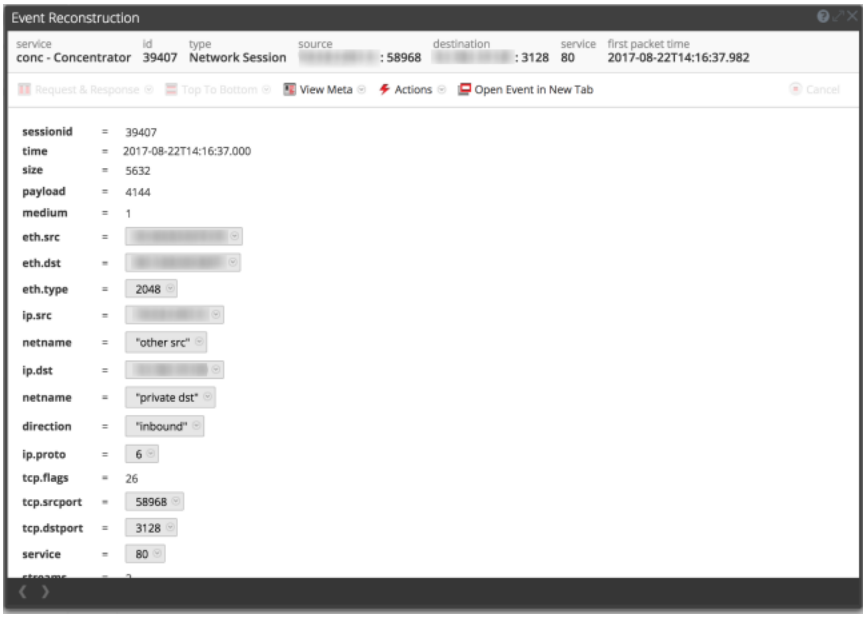
When you discover an event that merits additional investigation, you can analyze it in a form that best suits the content of event. Some forms of analysis involve safely reconstructing the event in a form similar to its native form, for example, packets, text, email, and web content. The rendering of events restricts the use of dynamic or active code that may be contained in the event to limit adverse outcome to your system or browser. Using cache improves performance when viewing previously viewed events. Each analyst has a separate cache of reconstruction data, and you can only access reconstructed events in your own cache.

Some network events are enriched with host data if you have an Endpoint deployment and the Endpoint Agent is configured for expanded network visibility. For such an event, you can view the host details.

The Events view gives you the ability to interactively analyze an event, looking at raw data, meta keys, meta values. This figure is an example of a network event rendered as packets in the Events view.



The Event Reconstruction in the Legacy Events view presents the raw data and the meta keys and meta values for an event in a list form. This figure is an example of the Event Reconstruction.



Configuring NetWitness Investigate Views and Preferences

Analysts can configure some aspects of NetWitness Investigate views and behavior. You can customize the way that Investigate views appear, the types of information displayed, and factors that affect performance in returning results and reconstructing events. All configurable settings have default values that are effective in most deployments; however, analysts have the option to adjust these if necessary.

Analysts using Investigate need to have the appropriate system roles and permissions set up for their user accounts. An administrator must configure roles and permissions as described in the *System Security and User Management Guide*.

These topics provide details:

- Configure the Navigate View and Legacy Events View
- Configure the Events View

Configure the Navigate View and Legacy Events View

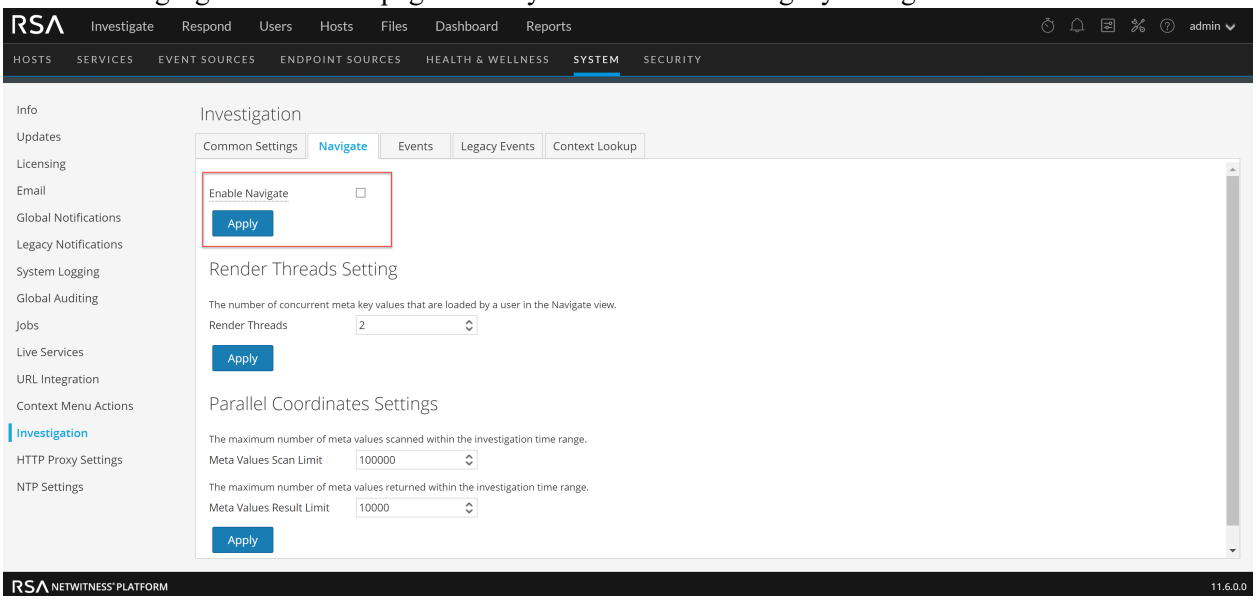
Analysts can set preferences that affect performance and behavior of NetWitness when using the Navigate view and Legacy Events view. Some of the same settings are available in two places in NetWitness, and changes made in either location are applied in the other view:

- Investigate view > Settings dialog for the Navigate view and the Legacy Events view.
- Profiles > Preferences panel > Investigation tab.
- Navigate view and Legacy Events view Search Options drop-down.

By default, the legacy Navigate view is disabled. To enable the Navigate tab in Investigate:

1. Go to  (Admin) > **System** > **Investigation** > **Navigate**.
2. Select the **Enable Navigate** checkbox.
3. Click **Apply**.

The following figure shows the page where you can enable the legacy Navigate view.



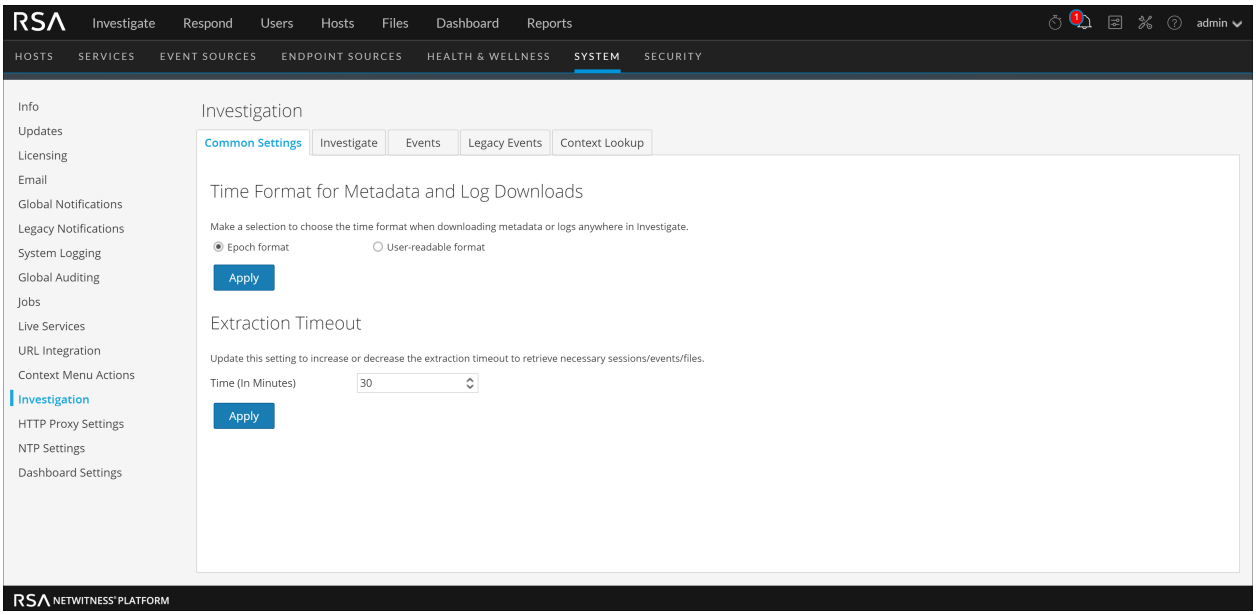
Configure Common Settings

In version 11.5 and later, the Common Settings tab allows you to configure settings that apply to the Navigate view, the Events view, and the Legacy Events view. You can set the time format used when downloading metadata and logs, and extraction timeout settings.

By default, the time format for downloads is Epoch format, which shows the time as a numerical value representing the number of seconds from the Unix epoch, January 1, 1970. The resulting number requires a conversion to be understood. You can change the setting to get a more understandable format that combines the user preference time zone, date format, and time format into an easily understood representation, which follows the industry standard ISO 8601 representation when possible.

This setting applies to all 11.5 Investigate views.

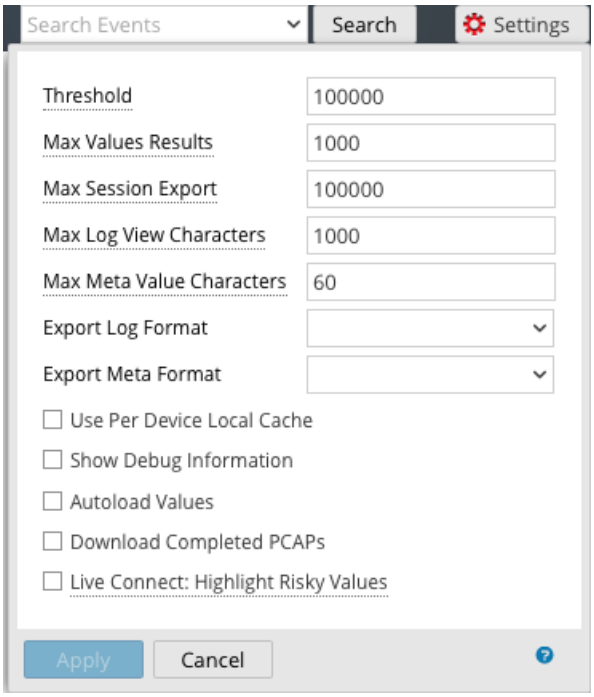
Go to **(Admin) > System**, and in the options panel, select **Investigation**.
The Investigation Configuration panel is displayed.



Access the Navigate View and Legacy Events View Settings

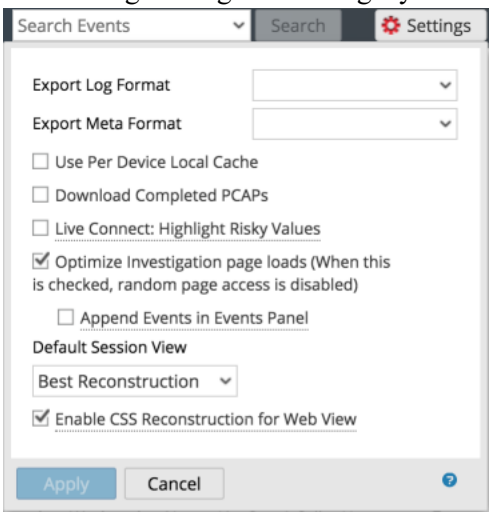
To access the settings, do one of the following:

- In the **Navigate** view toolbar, select the **Settings** option.
  The Settings dialog for the Navigate view is displayed.



Note: Version 11.0 included a setting to Append Events in Events Panel, which was moved to the Events view settings panel in Version 11.1.

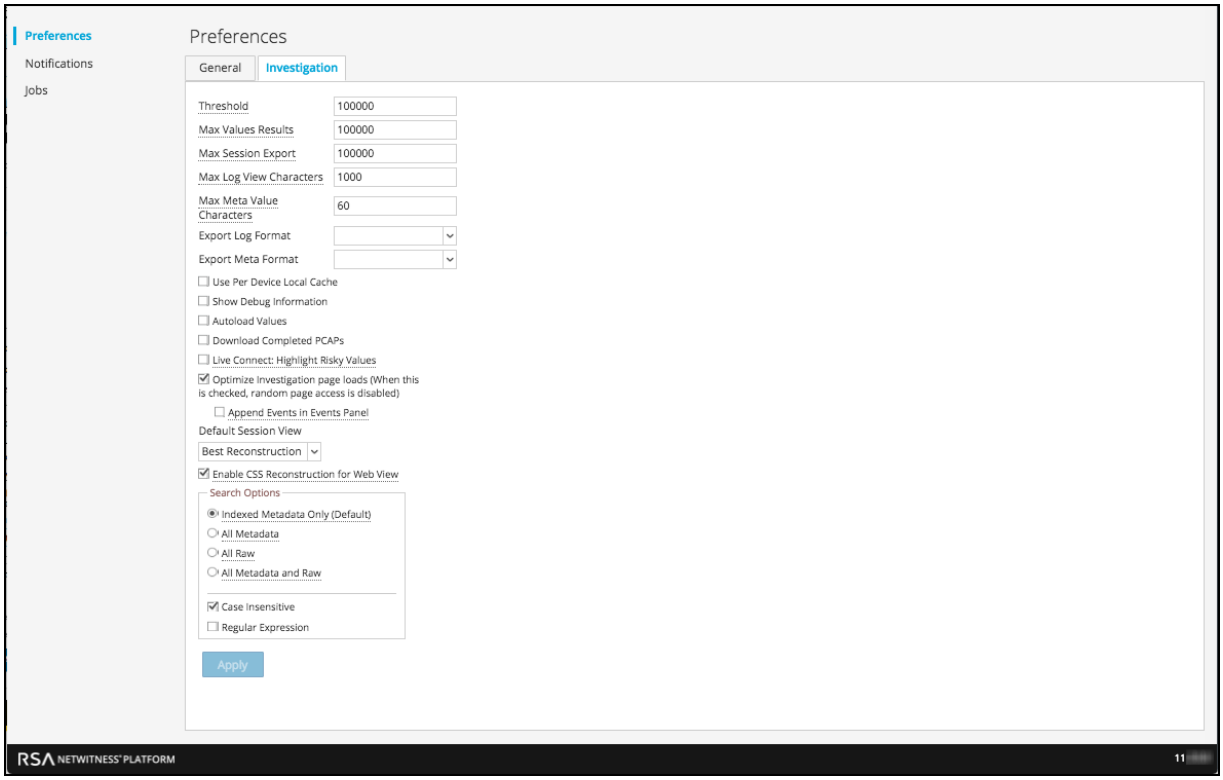
- In the **Legacy Events** view toolbar, select the **Settings** option.
  The Settings dialog for the Legacy Events view is displayed.



Note: Version 11.1 and later includes the Append Events in Events Panel setting.

- In the top right corner of NetWitness, go to > , Profile, and in the **Preferences** panel click the **Investigation** tab.

The Investigation panel is displayed. The figure below illustrates the Investigation panel.



Calibrate Navigate View Value Loading Parameters

Several settings influence the performance of NetWitness when loading values in the Values panel. Default values are set based on common usage, and individual analysts can adjust these settings for their own investigations. To adjust these settings:

1. Go to the **Preferences** panel > **Investigation** tab or to the **Settings** dialog for the Navigate view.
2. Adjust the following parameters:
   - **Threshold:** Set the threshold for the maximum number of sessions loaded for a meta key value in the Values panel. A higher threshold allows accurate counts for a value, and also causes longer load times. The default value is **100000**.
   - **Max Values Results:** Set the maximum number of values to load in the Navigate view when the Max Results option is selected in the Meta Key Menu for an open Meta Key. The default value is **1000**.
   - **Max Session Export:** Specify the number of events that can be exported in a single PCAP or Log file.
   - **Max Log View Characters:** Set the maximum number of characters to be displayed on **Investigate > Events > Log Text**. The default value is **1000**.
   - **Max Meta Value characters:** Set the maximum number of characters in a meta value name displayed in the Navigate view Values panel. The default value is **60**.

- **Show Debug Information:** If you want NetWitness to display the `where` clause beneath the breadcrumb in the Navigate view and the elapsed load time for each aggregated service on a Broker, check this option. The default value is **Off.**
- **Append Events in Events Panel:** This option affects paging in the Events view and is described below under "Calibrate Events View Retrieval and Default Reconstruction."
- **Autoload Values:** If you want NetWitness to automatically load values for the selected service in the Navigate view, check this option. When not selected, NetWitness displays a **Load Values** button, allowing the opportunity to modify options. The default value is **Off**.

3. Click **Apply**.
   The settings become effective immediately and are visible the next time you load values.

Configure Navigate View and Legacy Events View Parameters

Several settings influence the performance of NetWitness when loading values in the Navigate view and the Legacy Events view. Default values are set based on common usage, and individual analysts can adjust these settings for their own investigations. You can set these parameters separately in the Navigate view and the Legacy Events view. When configured in one view, the setting does not automatically apply to the other view. To adjust these settings:

1. Go to the **Preferences** panel > **Investigation** tab or to the **Settings** dialog for the Navigate view or the Legacy Events view.
2. Adjust the following parameters:
   - **Live Connect: Highlight Risky Values:** If you want NetWitness to highlight and display only IP addresses that are considered as risky by RSA community, check this option. When not selected, NetWitness displays all IP addresses. By default, this option is not selected (**Off**).
   - **Use Per Device Local Cache:** You can specify the use of locally cached data from the selected service. By default, this option is not selected **(Off)**. When unchecked, Investigate sends a fresh query to the database rather than displaying cached data in the Investigate views after the initial load. If checked, Investigate uses the data from local cache.
   - **Download Completed PCAPs:** You can automate the downloading of extracted PCAPs in the Navigate view and Legacy Events view so that the browser downloads the extracted PCAP and opens it in the default application for opening PCAP files, such as Wireshark. By default, this option is not selected (**Off**). If you are going to enable this option, ensure that an application that can open PCAPs is installed on your local file system and that the application is set as the default application to handle PCAP file formats.
   - **Live Connect: Highlight Risky Values:** If this option is unchecked, all the meta values that have context available in Live Connect are highlighted in the Navigate view Values panel. If the option is checked, among the values that have context in Live Connect, only those values deemed Risky/Suspicious/Unsafe by the community are highlighted. By default this option is unchecked (**Off**).
3. Click **Apply**.
   The settings become effective immediately.

Configure the Default Log Export Format

You can export logs from the Navigate view and the Legacy Events view as Text, XML, comma-separated values (CSV), and JSON. There is no built-in default value for the log export format. If you do not select a format here, NetWitness displays a selection dialog when you invoke export of logs. To select the format for exported logs:

1. Go to the **Preferences** panel > **Investigation** tab or to the **Settings** dialog for the Navigate view or Legacy Events view.
2. Select one of the options from the **Export Log Format** drop-down menu.
3. Click **Apply**.
   The setting goes into effect immediately.

Configure the Default Meta Value Export Format

You can export meta values from the Navigate view and Legacy Events view as Text, CSV, tab-separated values (TSV), and JSON. There is no built-in default value for the meta value export format. If you do not select a format here, NetWitness displays a selection dialog when you invoke export of meta values. To select the format for exported meta values:

1. Go to the **Preferences** panel > **Investigation** tab or to the **Settings** dialog for the Navigate view or Legacy Eventsview.
2. Select one of the options from the **Export Meta Format** drop-down menu.
3. Click **Apply.**
   The setting goes into effect immediately.

Note: If you upgrade to version 11.5.2, the **Export Meta Format** preference is not retained and is reset to blank. You must re-configure this value after you upgrade to version 11.5.2.

Calibrate Legacy Events View Retrieval and Default Reconstruction

You can configure several parameters that control the how NetWitness retrieves events and reconstructs events in the Legacy Events view. To adjust these paramaters:

1. Go to the **Preferences** panel > **Investigation** tab or to the **Settings** dialog for the Legacy Events view.
2. Configure the following parameters.
   - **Optimize Investigation page loads**: Set a paging option. When optimized, results are returned as quickly as possible,and you cannot go to a specific page in the event list. Unchecking this box changes the Events list pagination to allow you to go to a specific page in the list (or to the last page). The default value is **enabled**.
   - **Default Session View:** Selects the default reconstruction type for the initial reconstruction in the Legacy Events view. The default value is **Best Reconstruction** in which events are reconstructed using the reconstruction method most appropriate to the event.

3. Go to the **Preferences** panel > **Investigation** tab, or to the **Settings** dialog for the Navigate view (11.1) or the Legacy Events view (11.2 and later), and set the **Append Events in Events Panel** option. When this option is selected, the events displayed in the Events Panel are added incrementally. For example, each time you click the next page icon, the next increment of events is added, at first you see 1 to 25, then 1 to 50, then 1 to 75 and so on. This option is available only if the **Optimize Investigation Page Loads** option is enabled.
4. To activate the changes immediately, click **Apply**.

Enable or Disable Cascading Style Sheet Rendering in Web Content Reconstructions

Analysts can enable the use of cascading style sheets (CSS) when reconstructing web content. If enabled, the web reconstruction includes CSS styles and images so that its appearance matches the original view in a web browser. This includes scanning and reconstructing related events, and searching for style sheets and images used in the target event. The option is enabled by default. Disable this option if there are problems viewing specific websites.

Note: The appearance of the reconstructed content may not match the original web page perfectly if related images and style sheets could not be found or were loaded from the web browser's cache. Also, any layout or styling that is performed dynamically through the client side javascript is not rendered in the reconstruction because all client side javascript is removed for security purposes.

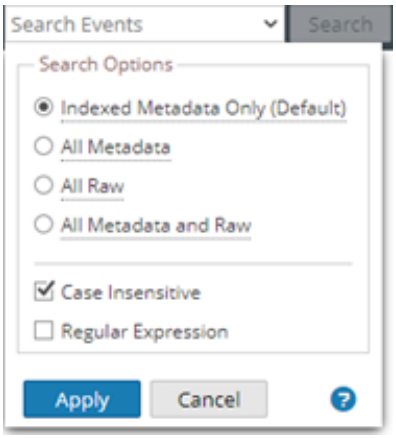
To enable or disable this option:

1. Go to the **Preferences** panel > **Investigation** tab.
2. Select the **Enable CSS Reconstruction for Web View** checkbox.
3. Click **Apply**.
   The setting becomes effective immediately and is visible in the next web content reconstruction.

Configure Search Options

You can configure search options to apply when you type a search string in the Search field. Edit the Search Options in the Profile > Preferences panel > Investigation tab or in the Navigate and Legacy Events view Search Options drop-down menu. To configure search options:

1. Go to the Search Options.
   The following figure illustrates the Search Options drop-down menu for Version 11.2 and later.



2. Select one or more search options to apply to the search. Search for Text Patterns in the Navigate and Legacy Events Views provides detailed information about each option.
3. To save the search settings, click **Apply**.
   The preferences are saved and effective immediately.

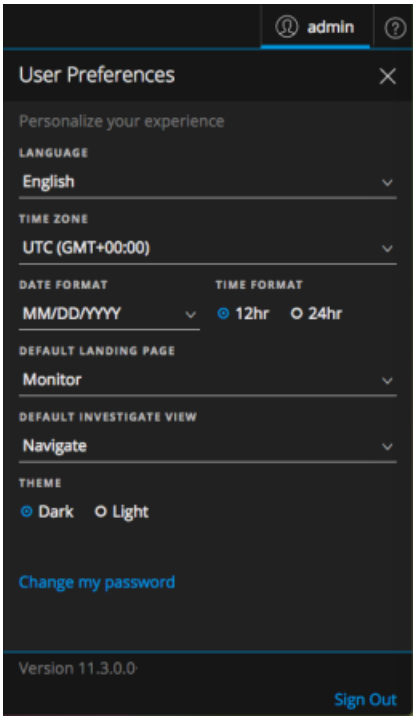
Configure the Events View

Note: The information in this topic applies to NetWitness Version 11.1 and later.

Analysts can set preferences that affect the behavior of NetWitness when using the Investigate > Events view. If the Events view is open; these two buttons give access to preferences dialogs: and . The User menu ( ) is focused on global user preferences such as time zone, while the Events preferences menu ( ) is focused on user preferences for behavior in the Events view. The rest of this section describes both sets of preferences.

Set the Default Investigate View

You can select the default view when you open Investigate: Navigate view, Events view, Hosts view, Files view, Entities view, or Malware Analysis view. The default Investigate view is set in the global User Preferences dialog (in the upper right corner of the NetWitness browser window, select). The global user preferences are described in detail in the *NetWitness Platform Getting Started Guide*.

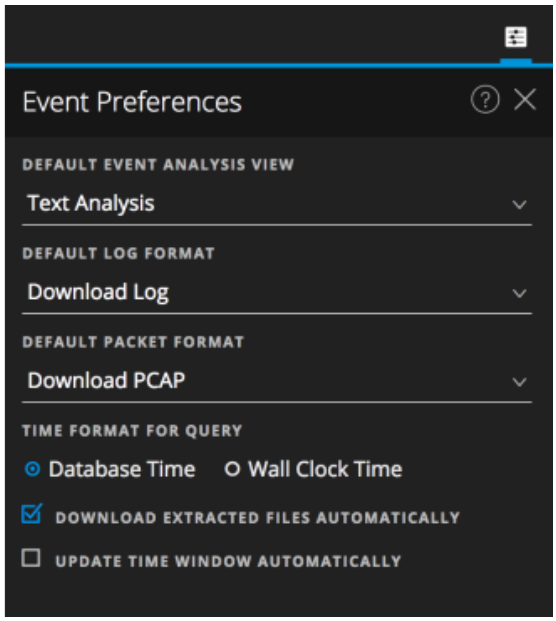


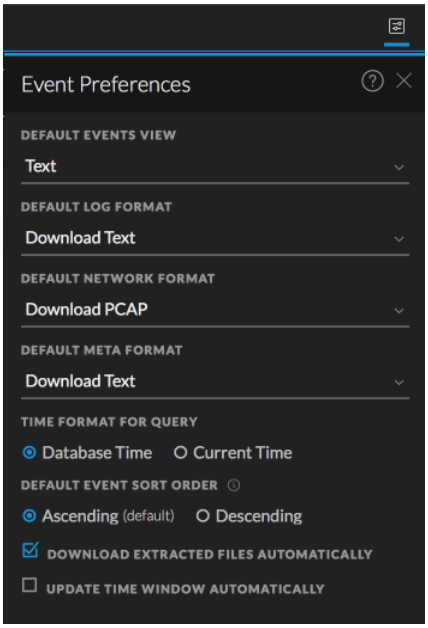
Set User Preferences for the Events View

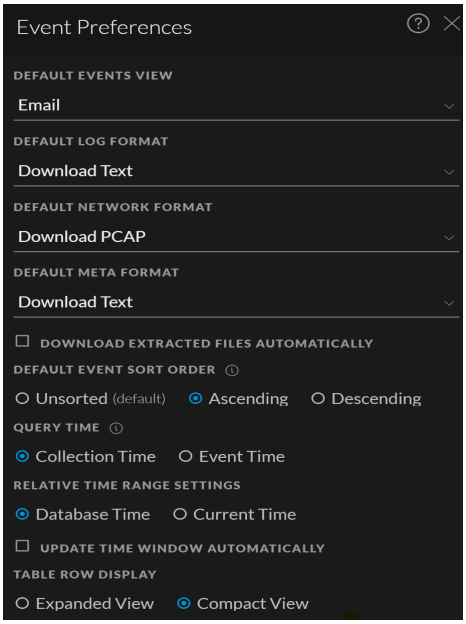
You can set your own preferences relevant to the Events view. The preferences selected persist per user and are available whenever the specific user logs in to the application.

To set default values for working in the Events view:

1. In the **Events** view, click .
   The Event Preferences dialog is displayed. Different versions of the dialog have some differences in labeling and available options as shown in the following figures. The first two figures are examples of the Version 11.3 and 11.4 dialogs respectively, and the third figure shows the dialog for Version 11.6.







2. In the **Default Events View** field, select the default reconstruction type when you open an event in the Events panel: **Text**, **Packet**, **File**, (Version 11.5 and Later) **Host** or **Email** .
   If you have not selected a default analysis type, when you open an event, the default reconstruction type is the Packet analysis, except for log and endpoint events, which open to the Text analysis. If you select a default reconstruction type, the reconstruction type is the default reconstruction that you specified. In both cases, the default is the starting point, and if you change the type while you are working, the type you choose is used for the next reconstruction.
3. In the **Default Log Format** field, select the download format for exporting logs: **Download Log** (11.3) or **Download Text** (11.4), **Download XML**, **Download CSV**, or **Download JSON**. If you do not select a format here, the default download format is **Download Text**. These options are also available at the time of download in a drop-down menu.
4. In the **Default Packet Format** (11.3) or **Default Network Format** (11.4) field, select the default format for downloading packets. If you do not select a format here, the default download format is **Download PCAP**. These options are also available at the time of download in a drop-down menu:
   - **Download PCAP** to download the entire event as a packet capture (*.pcap) file
   - **Download All Payloads** (11.3) or **Download Payloads** (11.4) to download the payload as a *.payload file
   - **Download Request Payload** to download the request payload as a *.payload1 file
   - **Download Response Payload** to download the response payload as a *.payload2 file
5. (Version 11.4 and later) In the **Default Meta Format** field, select the download format for exporting metadata: **Download Text**, **Download CSV**, **Download TSV**, or **Download JSON**. If you do not select a format here, the default download format is **Download Text**.
6. (Version 11.5.1 and later) To change the preference for the time that is matched when you submit queries, select an option in the **Query Time** field.
   - When **Collection Time** is selected, the time of an event reflects when it was received and stored into the system. This is the default setting.
   - When **Event Time** is selected, the time of an event reflects the time at which the event actually occurred. A good case for using the event time is when investigating logs or endpoints and looking for events that occurred around the same time. Using Event Time filters out all network events.

   Note: By default, the time range for a query is based on the time that the Decoder consumed the event, which is always not the same time that the event occured. To see the actual event time instead of the collection time in the Events view, select Event Time.

   Note: When you query with the event time preference enabled, you must use the collection time column and time zone column. This is to help you differentiate if the events are listed in a sequential order or not. This scenario occurs because there is no global standard to follow while logging the event times, which results in different events from different sources being in different time zones.
7. (Version 11.5.1 and later) Under **Relative Time Range Settings**, choose either **Database Time** or **Current Time**. The Events view can display results based on the relative time range, Last 2 hours or Last 30 Days, for example. The time range can be relative to the time when the event was received

and stored into the system or the current time zone's clock time. When you set the time format, your individual user preference is saved until changed again. The default setting for this preference is **Database Time**, which is the same time format used to display query results in the Navigate view and Legacy Events view.

- When **Database Time** is selected, the time range is relative to the time of the last stored event.
- When **Current Time** (labeled **Wall Clock Time** in Version 11.3 and earlier) is selected, the time range is relative to the current time in the timezone set in user preferences.

Note: (Version 11.6) **Current Time** is the default for **Relative Time Range Settings**. In previous versions, **Database Time** was the default value. Make a note that this may cause time range mismatch between **Events View** (using **Current Time** as default) and **Navigate View** (using **Database Time** as default). This change does not affect the existing users and is applicable only to the new users.

8. (Version 11.4 and 11.5) Under **Time Format for Query**, choose either **Database Time** or **Current Time**. The Events view can display results based on the database time or the current clock time. When you set the time format, your individual user preference is saved until changed again. The default setting for this preference is **Database Time**, which is the same time format used to display query results in the Navigate view and Legacy Events view.
   - When **Database Time** is selected, the end time for a query is based on the time that the event was stored.
   - When **Current Time** (labeled **Wall Clock Time** in Version 11.3 and earlier) is selected, the query is executed with the current time in accordance with the timezone set in user preferences.
9. (Version 11.4 and later) To set the sort sequence by collection time for the events listed in the Events panel, select one option under **Default Event Sort Order**. After you have selected a preference, you can also interact with the table column headers to request the results again sorted differently as described in Use Columns and Column Groups in the Events List.
   - **Unsorted** (default for Version 11.4.1): To list events as processed by the Core services. Unsorted is faster because it streams back the events as soon as a match is found versus waiting for all Core services to respond and then displaying them in the chosen order.
   - **Ascending** (default for Version 11.4 and earlier): To put the events with the earliest collection time first in the list. The oldest events are displayed first if in ascending order.
   - **Descending**: To put the events with the latest collection time first in the list. The newest events are displayed first if in descending order. When investigating logs, you may want to change the sort sequence to see the latest collection time first.
     If results exceed the events limit, not all events can be loaded. The portion of returned events loaded in the Events panel matches the sort order preference: the oldest portion of events is loaded when Ascending order is selected, and newest portion of events is loaded when Descending order is selected. When Unsorted is selected, the oldest portion of events is matched and then listed unsorted. if you changed the sort order preference after events were loaded, you must refresh the view to apply the new sort order.
10. If you want all extracted files to be downloaded automatically, select the **Download extracted files automatically** checkbox. You can go to the Jobs queue to view the extracted files.

11. (Version 11.3 and later) To automatically update the time range window in the query bar when the service is polled (at one minute intervals) and sends fresh results, select the **Update Time Window Automatically** checkbox. When the time range is updated, the  (Submit Query) button is activated and you can submit a query to get fresh results. To keep the time range window in the query bar synchronized with the current results, clear the checkbox (this is the default value
12. (Version 11.6 and later) If you have multiple rows of data and want to word wrap or unwrap the content of an event , choose either **Compact View** or **Expanded View**. Based on the selection you make the rows will appear compact or expanded.

Beginning an Investigation

Based on the question you are attempting to answer, NetWitness Investigate offers different starting points: Navigate view (Version 11.5 and earlier), Events view, Legacy Events view (Version 11.3 and earlier), Hosts view, Files view, Users (Entities) view, and Malware Analysis view.

Specific user roles and permissions are required for a user to conduct investigations in NetWitness. If you cannot perform a task or see a view, the administrator may need to adjust the roles and permissions configured for you.

Note:
- The Files and Hosts views are available in Version 11.1 and later (refer to the *NetWitness Endpoint Quick Start Guide* and *NetWitness Endpoint User Guide* for details). Before Version 11.5, these views were submenu of Investigate.
- The Users view is available in Version 11.2 and later (refer to the *NetWitness UEBA Quick Start Guide* and the *NetWitness UEBA User Guide* for details); in Version 11.4 it is labeled Entities view. Before Version 11.5, it was submenu of Investigate.
- By default, the Legacy Events view is disabled in Version 11.4, but can be enabled by an administrator as described in the *System Configuration Guide*.
- By default, the Navigate view is disabled in Version 11.6 as the Filter Events Panel in the Events view provides this functionality. To enable the Navigate view, see Configure the Navigate View and Legacy Events View.
- Specific user roles and permissions are required for a user to conduct investigations and malware analysis in NetWitness. If you cannot see a view, the administrator may need to adjust the roles and permissions configured for you.
- The 11.4 Events view is the default view for investigating events. The default workflow for analysts interacting with events is optimized to limit the need to transition from one view to another. By combining capabilities that were previously in two distinct workflows, known as Event Analysis and Events, the analyst has a single workflow for analyzing events. With the new functionality added to the Events view, the Legacy Events view is no longer needed. By default the previous workflow is not in the Investigate menu, but an administrator can re-enable it as described in "Configure Investigation Settings" in the *System Configuration Guide*.

Focus on Metadata, Raw Events, and Event Analysis

To hunt for events that drive the incident response workflow and to do strategic analysis after another tool has generated an event, go to **Investigate > Navigate**, **Investigate > Events**, or **Investigate > Legacy Events**. You can investigate the metadata and raw events for a single Broker or Concentrator. In each of these views, you can execute a query and filter the results by narrowing the time range and querying metadata. These topics provide details about beginning an investigation:

- Begin an Investigation in the Events View
- Begin an Investigation in the Navigate or Legacy Events View

Focus on Hosts and Files

To hunt for information on hosts that have the Endpoint agent running, go to **Hosts** (Version 11.5) or **Investigate > Hosts** (Version 11.4). For every host, you can see processes, drivers, DLLs, files (executables), services, and autoruns that are running, and information related to logged-in users. To begin an investigation by looking at files in your deployment, go to **Investigate > Files**. (See the *NetWitness Endpoint User Guide* for detailed information.)

Focus on Risky User and Entity Behavior

To discover, investigate, and monitor risky behaviors across all users and entities in your network environment, go to **Users** (Version 11.5), **Investigate > Entities** (Version 11.4), or NetWitness UEBA (User and Entity Behavior Analytics) . In versions 11.3 and earlier, the menu option is **Investigate > Users**. You can detect malicious and rogue users, pinpoint, high-risk behaviors, discover attacks, and investigate emerging security threats. (See the *NetWitness UEBA User Guide* for detailed information.)

Focus on Scanning Files for Malware

To scan files for potential malware, or set up a continuous scan of a service, go to **Investigate > Malware Analysis**. Scan results are expressed as four types of analysis: network, static, community, and sandbox with an indicator of compromise (IOC) rating. There are several other ways to begin working in Malware Analysis:

- You can begin Malware Analysis from the Malware Analysis dashlets in the Monitor view to quickly see the riskiest potential threats.
- You can right-click a meta key in the Navigate view, and select **Scan for Malware**.

See the *Malware Analysis User Guide* for detailed information.

Begin an Investigation in the Navigate or Legacy Events View

The Navigate view is the default view for Investigate unless you have selected a different view as your opening view. This user preference is set on the application level as described in Configuring NetWitness Investigate Views and Preferences. In the Navigate view and Legacy Events, you are hunting for events of interest based on a query. In the Navigate view you can also refine results by clicking on meta keys and meta values. When you find interesting events, you can take a closer look at the event in the other Investigate views.

To begin an investigation in the Navigate view or Legacy Events view, a service must be specified.

- Investigate opens the Navigate view or the Legacy Events view with the user-specified default service selected.
- If no default service is currently specified and the service id is not in the URL, Investigate presents a dialog for selecting the service or collection to investigate.
- When a service is selected manually or by default in the Navigate view or Legacy Events view, you can change the service or collection to investigate by selecting the service name in the toolbar. Investigate presents the dialog for selecting the service to investigate.

Note: The Archiver service does not appear in the Navigate view to minimize user experience of slow performance when performing investigations. The Archiver is available in the Legacy Events view for log exports and enhanced search capabilities.

With a service or collection selected, Investigate is ready to load data for the service or collection. It is recommended that you also select a time range so that results load faster. Several settings in the Navigate view and Legacy Events view Settings dialog or the Profiles > Preferences panel > Investigations tab affect the loading process: Threshold, Max Values Results, Show Debug Information, Autoload Values, and Optimize Investigation page loads (see Configuring NetWitness Investigate Views and Preferences).

Note: In the Legacy Events view data loads automatically. If you specified Autoload Values in the Navigate view preferences, Investigate populates the data automatically. Otherwise, you must select Load Values. Investigate populates the metadata in the Navigate view Values panel and results become visible almost immediately.

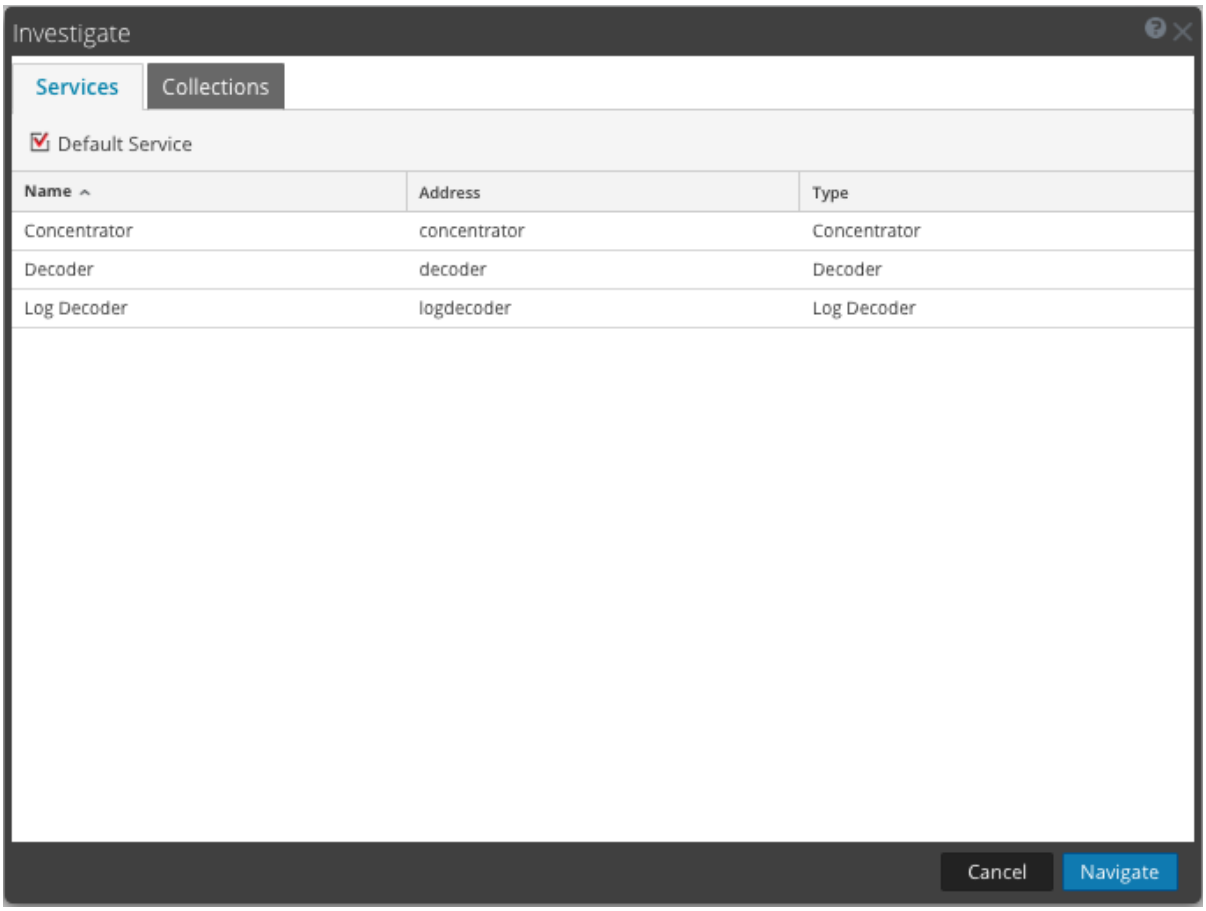
The rest of this topic provides instructions for beginning the investigation of data on a service.

Note: Only users with the administrator role can create a collection, and only the creator of the collection is able to investigate a collection.

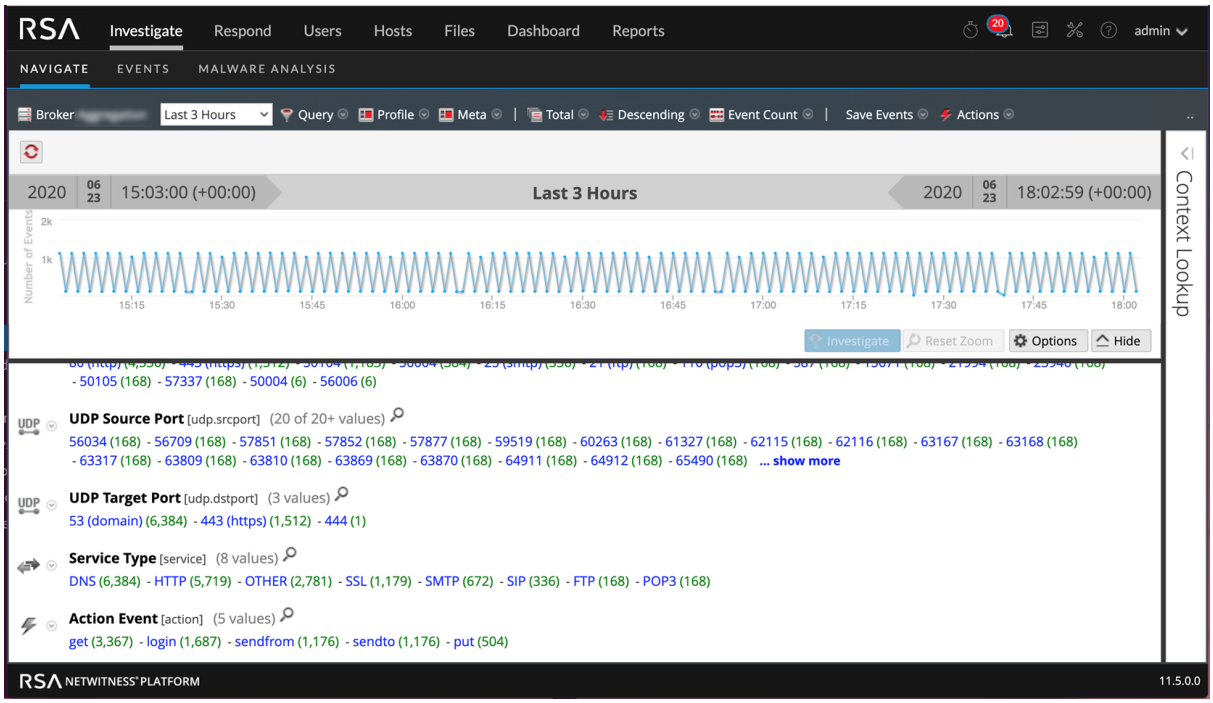
After loading data in the Navigate or the Legacy Events view, refine results, reconstruct and analyze events, then download and act upon results (see Refining the Results Set and Reconstructing and Analyzing Events and Downloading and Acting Upon Results).

Begin an Investigation (No Default Service)

1. Go to **Investigate > Navigate** or **Legacy Events**.
   The Investigate dialog is displayed.



2. Double-click a service or select a service, usually a Concentrator, and click **Navigate**.
   The data loads automatically in the Legacy Events view. If you are working in the Navigate view, the resulting panel displays the activity for the selected service, but the data is not loaded automatically.
3. (Recommended) Select a specific time range so that results load faster.
4. If you want to modify investigation options before loading, you can create or modify a custom profile, apply a different time range, create or apply a meta group, and perform a custom query as described Refining the Results Set. You can also modify options at any time during the investigation.
5. To load data in the Navigate view, click Load Values.
   The data for the selected service begins loading.

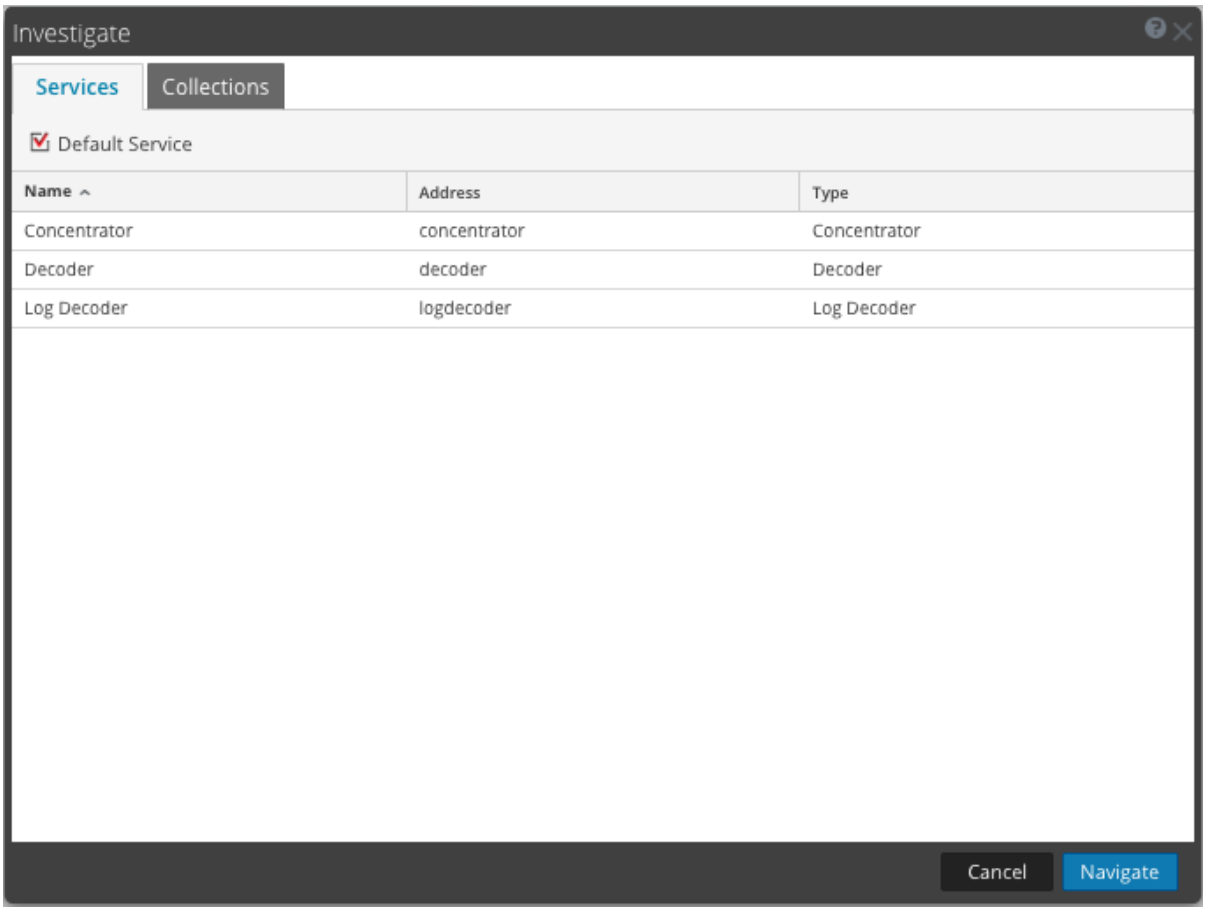


With the service selected and data loaded, you are ready to begin analyzing the data.

Set or Clear the Default Service

You can set the default service and clear the default service in the Investigate a Service dialog.

1. Click the service name in the toolbar.
   The Investigate dialog is displayed.



2. Select a service on the **Services** grid, and click Default Service.
   The service becomes the default, (indicated by **Default** in parentheses after the service name).

3. To clear the default service, select the default service in the grid, click Default Service, and click **Cancel** to close the dialog.
   No default service is set.

> **Note:** Clicking cancel does not cancel your selection of the default service. It closes the dialog without navigating to the currently selected service in the grid. Setting a default service that is different from the service currently being investigated, does not refresh the Navigate view. You must explicitly select and navigate to a different service.

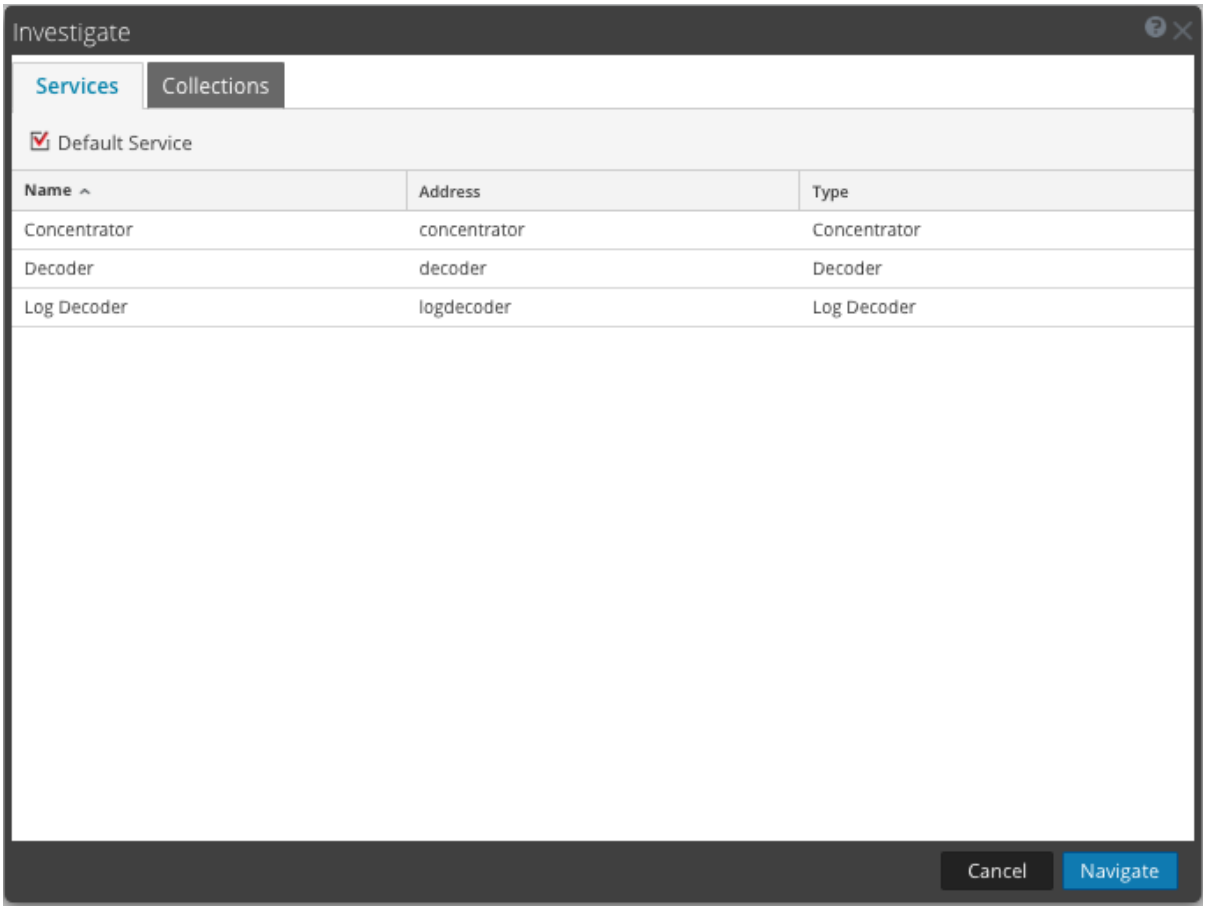
Begin an Investigation (Default Service Specified)

1. Go to **Investigate > Navigate** or **Legacy Events**.
   If the Autoload Values setting is set to off, the Navigate view is displayed with the default service selected, and ready to load data. If the Autoload Values setting is on, the values are loaded as shown in Step 3. In the Legacy Events view, the data is loaded automatically.

2. If you want to modify investigation options in the Navigate view before loading, you can create or modify a custom profile, apply a different time range, create or apply a meta group, and perform a custom query.
3. When ready, click Load Values.
   The values for the service are loaded in accordance with the selected options. With the service selected and data loaded you are ready to begin analyzing the data.

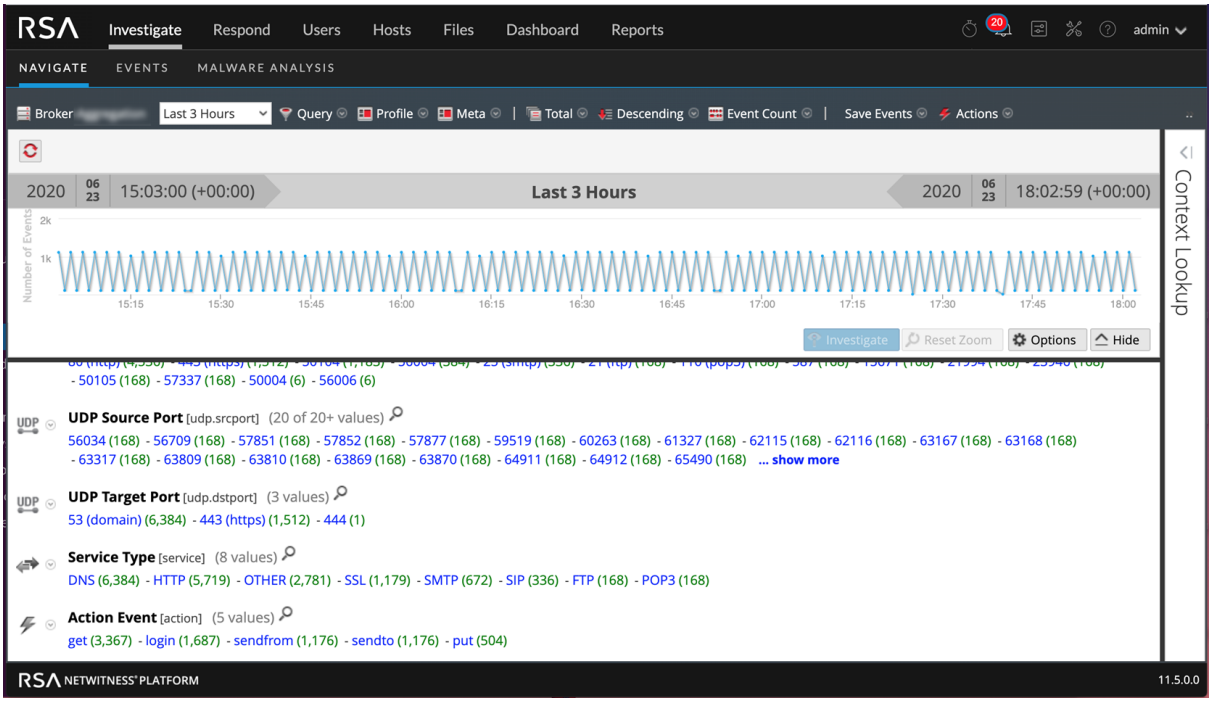
Change the Service or Collection to Investigate

1. In the Navigate view or the Legacy Events view, click the service name at the top of the options panel.
   The Investigate dialog is displayed.



2. Double-click a service or select a service and click **Navigate**. The resulting panel displays the activity for the selected service.
   If the Autoload Values setting is on, the values are loaded as shown in Step 3. Otherwise, the Navigate view is displayed with the default service selected, and data ready to load. In the Events view the data is loaded automatically.

3. When ready, click Load Values.
   The values for the service begin loading in accordance with the selected options.



   With the service selected and data loaded you are ready to begin analyzing the data.

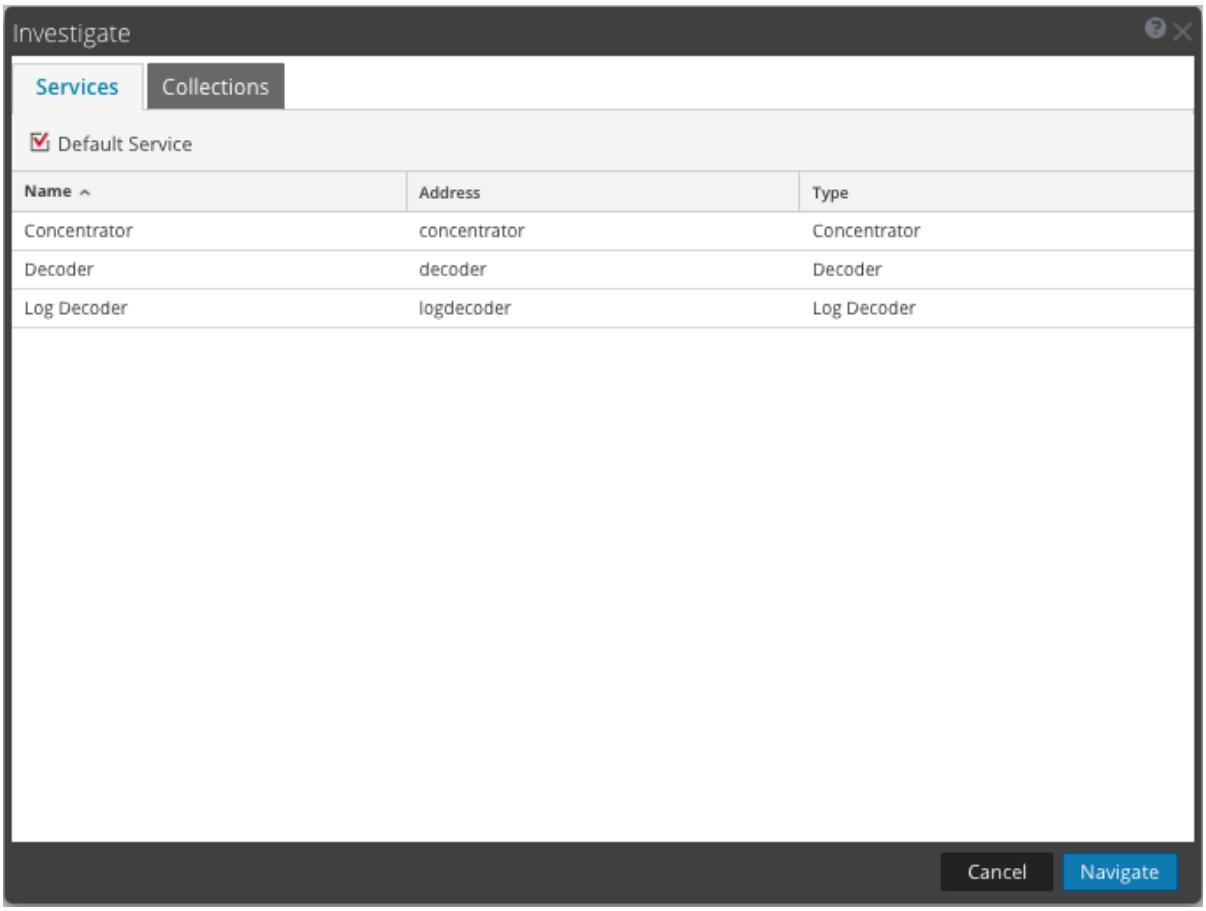
Investigate Workbench Restoration Collections

This procedure enables administrators to select content from an existing collection to reprocess for further investigation. This applies to Decoders that use Workbench services.

Note: Only a user with administrative privileges can create a collection, and you can view only those collections that you created.

To reprocess data for further investigation:

1. Go to **Investigate > Navigate** or **Legacy Events**.
   The Investigate dialog is displayed.

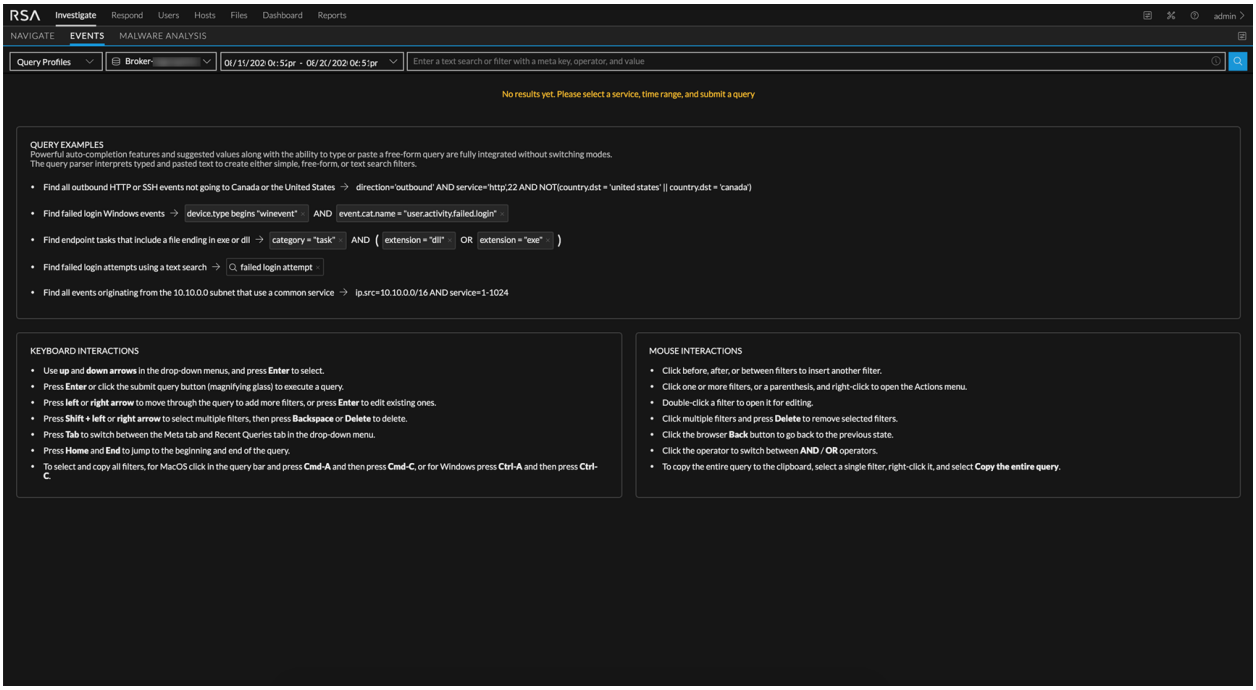


2. Select a workbench service and workbench name that you want to investigate.
3. Click **Navigate** to perform an investigation on the selected workbench service.
   Click **Cancel** to select a different workbench service to investigate.
   The Investigation view is displayed. With the collection selected and data loaded, you are ready to begin analyzing the data.

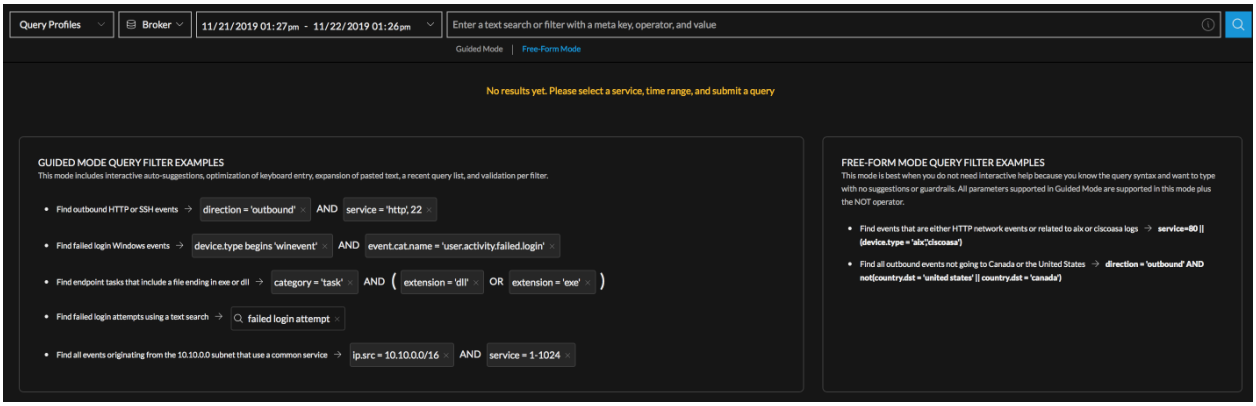
Begin an Investigation in the Events View

The Events view offers most of the features that are available in both the Navigate view and the Legacy Events view. Similar to the Navigate view, there is a view into meta keys and meta values for logs, endpoints, and packets. Like the Legacy Events view, an events list shows events listed in the order by time, and you can view the raw event, related metadata, and a reconstruction of an event. The Event reconstruction has some helpful cues to identify points of interest. See Reconstructing and Analyzing Events.

The following figure shows the initial Events view with some examples of queries and information about keyboard and mouse interaction. This figure depicts the initial view.



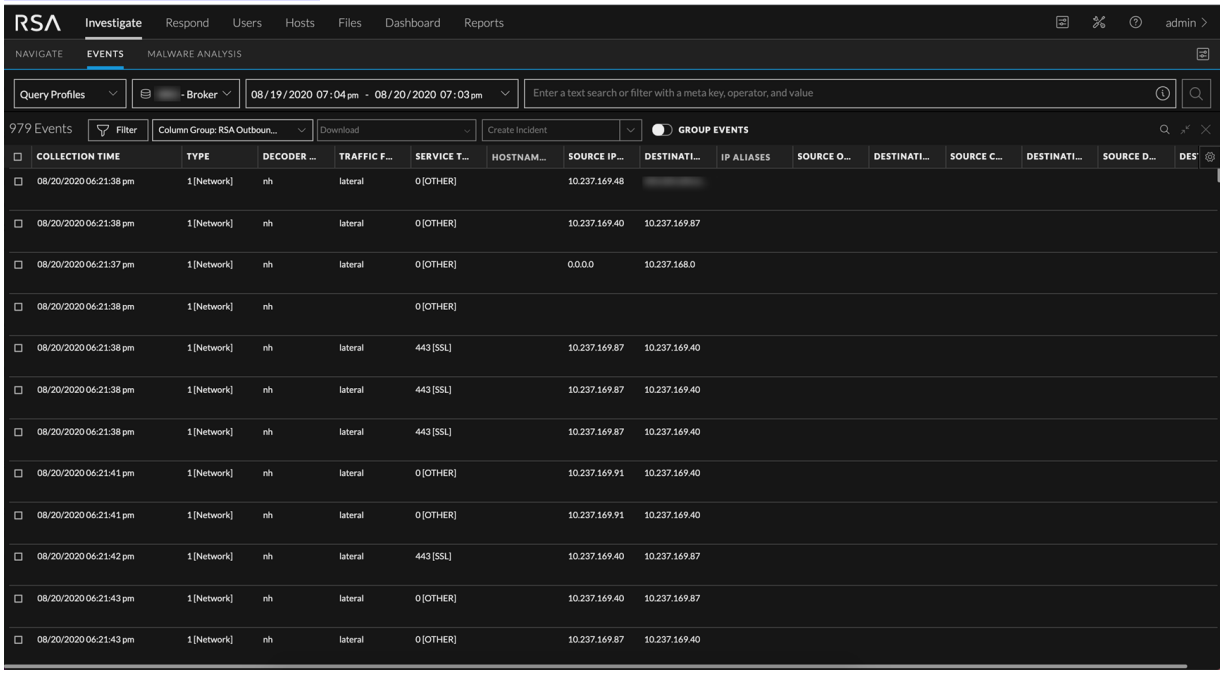
This figure show the examples provided in Version 11.4, which has two modes for creating queries.



Access the Events View

Several ways to access the Events view are available in Version 11.1 and later.

- Go to **Investigate > Events** or select the **Investigate** option in the main menu if you have made the Events view your default Investigate view. The following procedure provides detailed steps.
- Hover over and click a count (the green number after a meta value) in the Navigate view. The Events view opens with the list of events for the selected drill point, and you can begin working as described in Analyze Events in the Events View.
- Hover over a count and control-click **Open Events in new tab**. The Events view opens in a new tab with the list of events for the selected drill point, and you can begin working as described in Analyze Events in the Events View. The following figure is an example of the list of events.

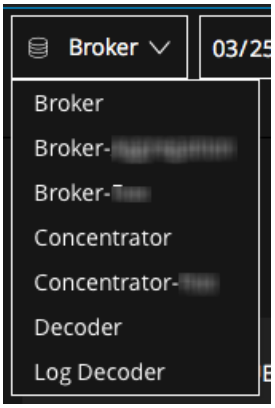


To begin an investigation in the Events view using direct access:

1. Go to **Investigate** > **Events**.
   The Events view opens with a service selected and no data displayed. A drop-down list offers a list of available services in alphabetical order. The **Select a service** field is populated with the first service in the list or the most recently selected service. By default the list of available services is retrieved every twelve hours and cached on the NetWitness server. If a service is added or removed from the NetWitness server before the next time to retrieve, the cache is updated with the latest list of services. An icon provides the status of the service.

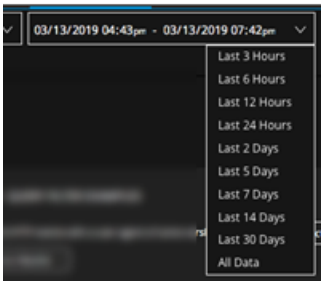
- and selected service name = The service is selected.
- = Investigate is attempting to connect to the selected service.
- = There was an error connecting to the selected service or there is no data in the selected service. In this state, the service selector control also turns red, and a tooltip explains why the connection attempt failed and advises you to choose another service.

2. (Optional) Select a service, usually a Broker or Concentrator, from the drop-down list.



The time range selector shows either the default time range of 24 hours, or the time range that you last selected for this service. The or Query Events button becomes active and you can create filters. If you launch a query without creating filters, the selected time is used.

3. (Optional) Edit the time range as described in Filter Results in the Events View.



The selected time range is stored in your browser for this service; you can set different time ranges for different services.

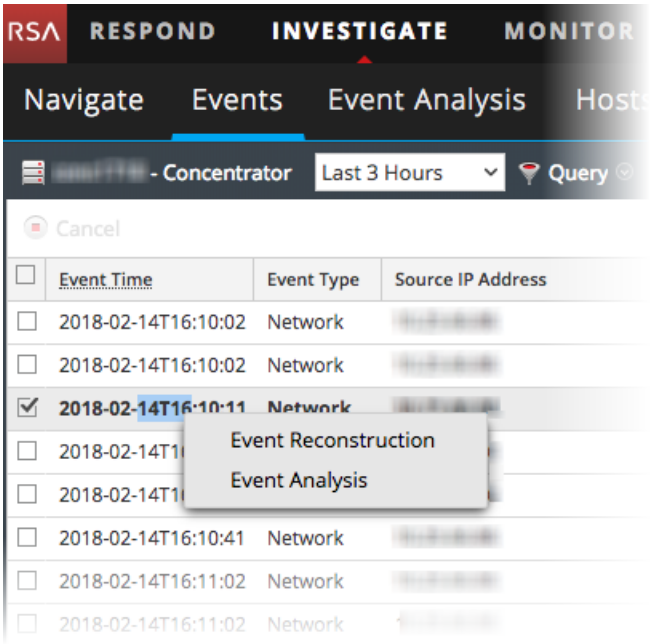
4. Create a query that consists of one or more filters that contain a meta key, operator, and optional value. See Filter Results in the Events View for details on creating queries.

5. When ready to submit the query, click or **Query Events**.
The Events view displays the data for the selected service, time range, and query, in accordance with permissions assigned to your role by the administrator. You are ready to begin analyzing the data. Refer to Examine Event Details in the Events View and Analyze Events in the Events View to learn how to work in the Events view.

Access the Events View (Version 11.0)

To open an event in the Event view:

1. Go to **Investigate** > **Events**.
2. Do one of the following:
   a. Right-click an event in the listed events, and select **Event Analysis**.



   b. Right-click an event in the listed events and select **Event Reconstruction**. Then click the **Event Analysis** button in the reconstruction.

Refer to Examine Event Details in the Events View and Analyze Events in the Events View to learn how to work in the Events view.

Refining the Results Set

When conducting an investigation, results load faster and it is easier to find what you are looking for if you refine the results to get a smaller number of results. In addition, limiting the time range and submitting a good query gives you more relevant results to answer the question at hand. Use a combination of the methods described in the rest of this section to get the information you need quickly.

- Use Meta Groups to Focus on Relevant Meta Keys
- Use Columns and Column Groups in the Events List
- Use Query Profiles to Encapsulate Common Areas for Investigation
- Filter Results in the Events View
- Filter Results in the Navigate View
- Filter Results in the Legacy Events View
- Create a Query in the Navigate and Legacy Events Views
- View and Modify Queries Using URL Integration
- Search for Text Patterns in the Navigate and Legacy Events Views

Use Meta Groups to Focus on Relevant Meta Keys

A meta group combines selected meta keys and meta entities into a group to show only data in which the meta keys and meta entities were found. In the Navigate view and the Version 11.5 and later Events view, you can use meta groups to filter data displayed in the Navigate view (Values panel) and the Events view (Filter Events panel). The same shared meta groups are available for use in both views. Private meta groups created in the Events view are not available for use in the Navigate view or in query profiles in the Legacy Events view.

Note: In the Navigate view and Legacy Events view, you can manually add non-indexed meta keys (or keys that are not in the index at all) to a meta group or column group. The non-indexed meta keys are fully available (manageable and displayable) in the Navigate view and Legacy Events view, but only partially (displayable in the Filter Events panel) in the Events view. The Events view (Filter Events panel) can display data for non-indexed meta keys that are already included in a meta group, but you cannot add non-indexed meta keys while you are editing a meta group. The non-indexed meta keys in a column group do not display data in a column and new non-indexed meta keys cannot be added to a column group in Events view.

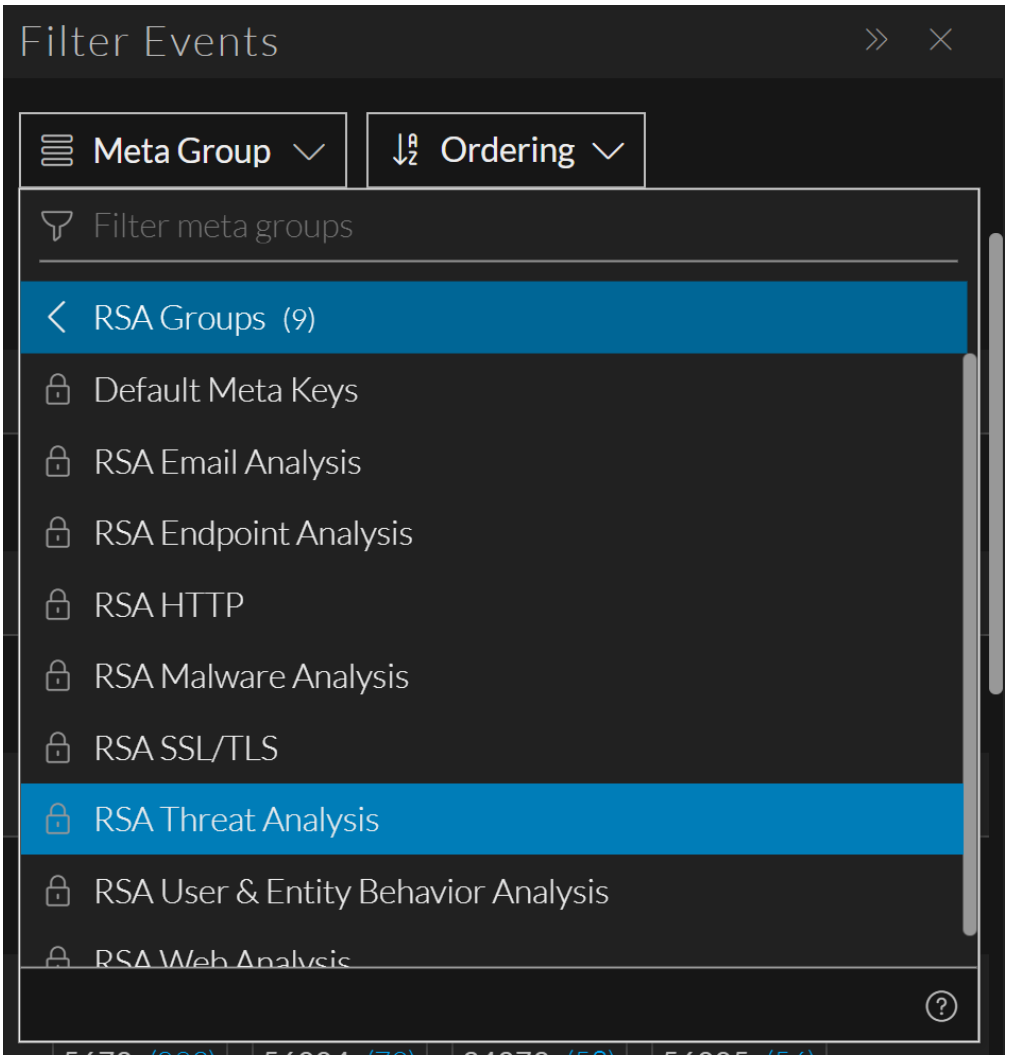
With a meta group in effect during an investigation, the information in the Values panel or the Filter Events panel shows only the meta keys in the selected group. When you open a Parallel Coordinates visualization in the Navigate view, the meta keys and meta entities in a group appear as axes from left to right. It may be useful to create two versions of each custom meta group; one for analysis of meta values and one for creating a parallel coordinates chart focusing on a smaller subset of the same use case.

A fresh installation of NetWitness includes built-in meta groups to help you find interesting data sets in Investigate. The built-in meta groups can be duplicated but cannot be edited or deleted. You can also create your own groups and edit a copy of a built-in group to create a custom group.

All groups in the Navigate view are shared and visible to all users of a service; you can export a group for import to any service, limited by the available meta keys for that service. In the Version 11.5 Events view Filter Events panel, you can create both shared and private custom meta groups; only the shared groups are visible and usable in the Navigate view.

Live Meta Groups

In 11.6 and later, NetWitness supports deploying the investigate content from live. The meta groups are categorized as RSA Groups (RSA Live content and RSA OOTB Groups), and Shared Groups. The content deployed from Live are marked by the live symbol (). The content is displayed in a folder structure. The groups are displayed as non-editable folders and sub-folders. The number inside **()** depicts the number of contents inside a folder and **>** symbol helps you to drill down inside the folder.



Built-In Meta Groups

RSA NetWitness has built-in meta groups, prefixed with RSA, that are available immediately after installation. The built-in meta groups are useful to focus an investigation on common use cases and to support threat detection using the RSA Hunting Pack. You can copy these groups, give the copy a new name, then edit the copy. These are the built-in meta groups:

- RSA Email Analysis includes meta keys that outline email interactions.
- RSA Endpoint Analysis contains meta keys that provide insight on processes, files, users, and connections from NetWitness Endpoint (NWE) hosts.
- RSA Malware Analysis includes meta keys that mark indicators of compromise in files contained in events.
- RSA HTTP includes meta keys that provide insight into outbound web traffic.
- RSA SSL/TLS includes meta keys that focus on encrypted web traffic.
- RSA Threat Analysis includes meta keys that mark potential threats in the data set.

- RSA User & Entity Behavior Analysis includes meta keys that encompass all the meta keys to analyze user and entity behavior.
- RSA Web Analysis includes meta keys that mark anomalies in web traffic.

Default Meta Keys Group (Version 11.5 Events View)

The Default Meta Keys meta group is a special type of built-in meta group that consists of all the meta keys for the currently selected service, returned in the order of appearance in the index file for the service. The Default Meta Keys group is always at the top of the list in the Meta Groups menu.

The Default Meta Keys group is used to select meta keys shown in the Filter Events panel when no meta group has been selected and none exists in local storage. You can also select this group as you would any other group. When using the Default Meta Keys group in the Filter Events panel, only the first 30 meta keys with values are open and the remaining are closed.

Custom Meta Groups

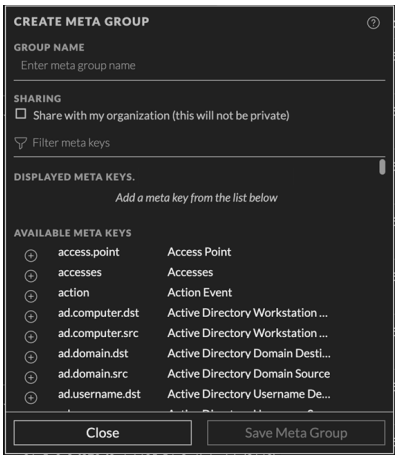
You can create custom meta groups to support scenarios that you use frequently while working in Investigate. When an administrator adds custom meta groups manually by editing the custom index file for a service, the new meta groups become available to use in meta groups after the service is restarted. Custom meta groups can be shared or private . Shared meta groups are available globally within your organization in the Navigate view and in the Filter Events panel. If you edit a shared custom meta group, your changes are applied globally. If you delete a shared custom meta group, the group is deleted and no longer available for all analysts. The Navigate view supports only shared groups. When you create custom meta group in the Events view, you can choose to share it or you can keep it private (default); you cannot change a shared group to private or a private group to shared.

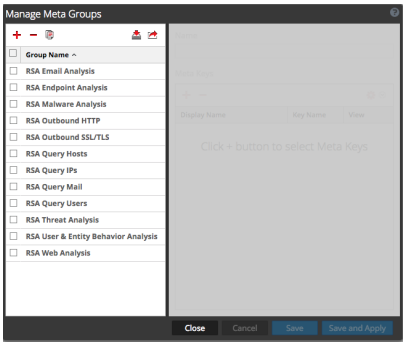
Note: Private custom meta groups created in the Events view are not visible or usable in the Navigate view.

Icons identify the group type in the Meta Groups menu. These are examples of each type of custom meta group with the edit icon displayed at the end of the row.



While the functionality of meta groups is similar in the Navigate view and the Events view, the user interface and some of the procedures are different. The following figures illustrate the (Events view) Create Meta Group dialog and the (Navigate view) Manage Meta Groups dialog.





Using options in the Events view Meta Groups menu (Version 11.5 and later) , you can:

- Select a meta group to apply.
- See the details of a meta group.
- Create, edit, and delete custom meta groups.
- Copy and edit the copy of a built-in or custom meta group.

Using options In the Navigate view Manage Meta Groups dialog, you can do all of the above as well as import and export a meta group.

The rest of this topic provides instructions for working with meta groups in the 11.5 Events view and the Navigate view.

Work with Meta Groups in the Events View (Version 11.5 and Later)

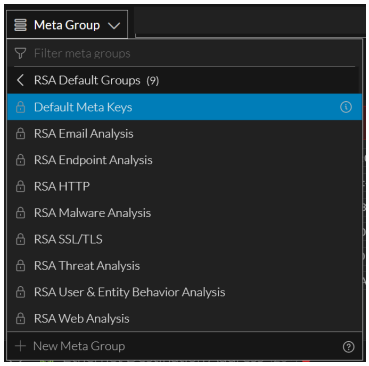
After the upgrade to Version 11.5 or later, all of the existing meta groups -- both built-in and custom -- are available for filtering events in the Filter Events panel. The meta group selection persists between logins unless browser cache is cleared.

View the Meta Keys in a Meta Group

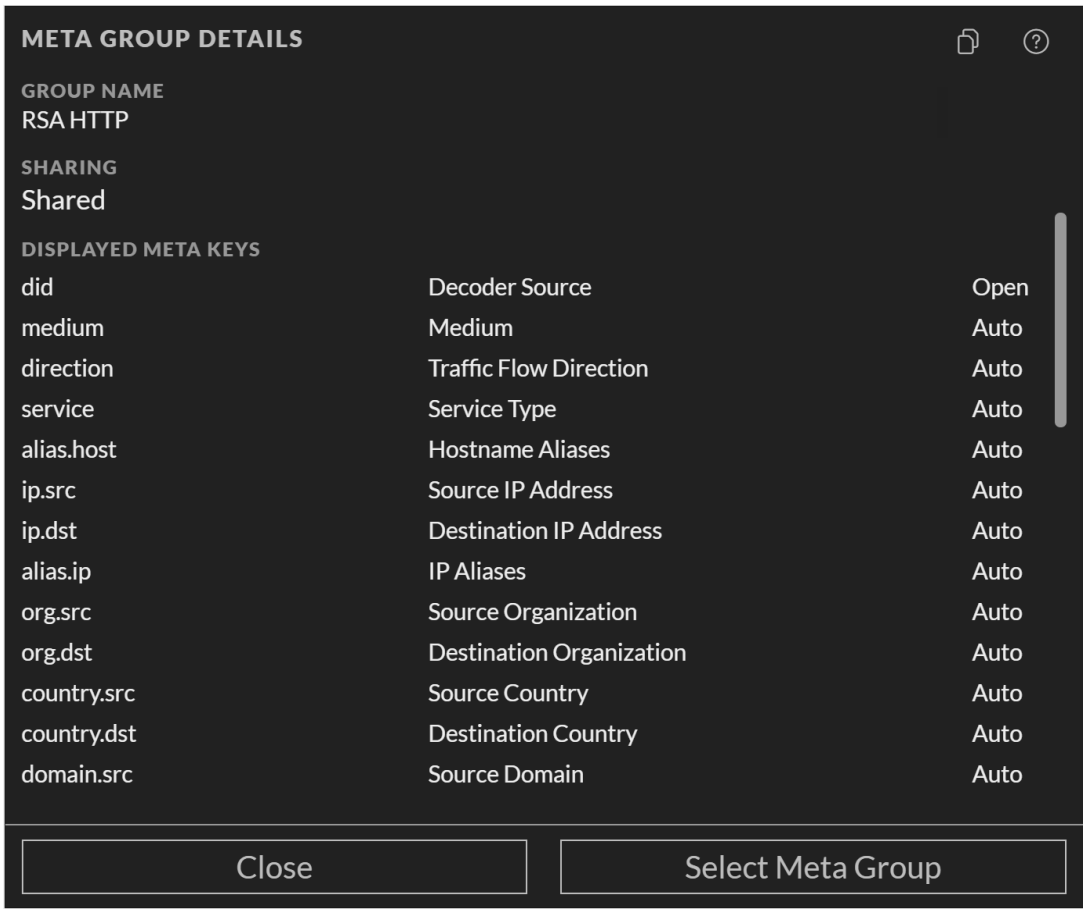
To view details of a meta group:

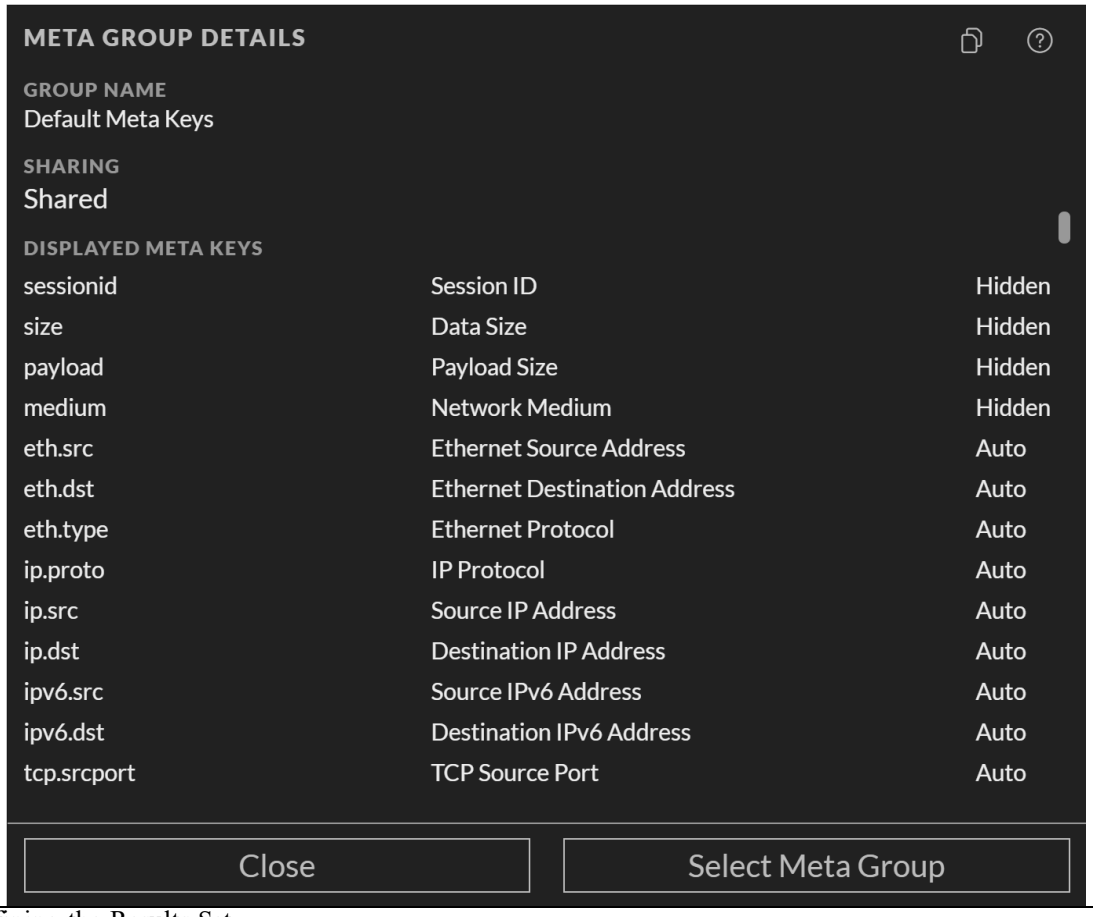
1. Go to **Investigate > Events** and click  to load events.
   The events for the default service and the default time range are loaded in the Events panel.

2. To display the Filter Events panel, click Filter above the Events panel.
   The Filter Events panel opens to the left of the Events panel.
3. To display the Meta Groups menu, click the Meta Groups menu title. The menu title is either Meta Group: Default Meta Keys or Meta Group: <currently selected meta group>. If this is your first visit after logging in, the Default Meta Key group is selected; any subsequent visits use the meta group selected in the previous session. If the selected meta group from the previous session is deleted, the Default Meta Keys group is selected when you log in. When opened, the menu displays a list of built-in meta groups (RSA), shared custom meta groups, and your private custom meta groups. Above the list, visibility options and a filter make it easier to find a particular meta group.



4. (Optional) To filter the listed meta groups by name, type some text in the **Filter Meta Groups** field. The list is updated to show only the group names that contain the exact text.
5. Hover over the meta group name and click the information icon ( ) to see which meta keys are included in the group.
   The figure on the left shows the columns for the RSA HTTP meta group. The figure on the right shows the columns for the Default Meta Keys meta group.

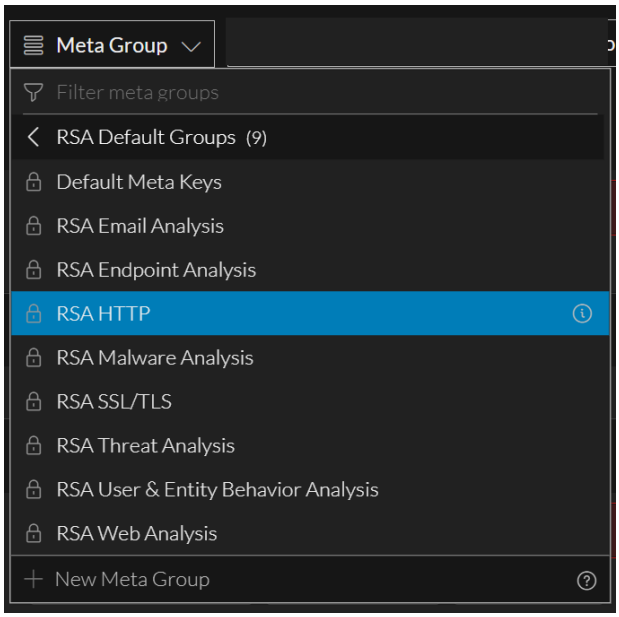

META GROUP DETAILS
GROUP NAME
RSA HTTP
SHARING
Shared
DISPLAYED META KEYS
did Decoder Source Open
medium Medium Auto
direction Traffic Flow Direction Auto
service Service Type Auto
alias.host Hostname Aliases Auto
ip.src Source IP Address Auto
ip.dst Destination IP Address Auto
alias.ip IP Aliases Auto
org.src Source Organization Auto
org.dst Destination Organization Auto
country.src Source Country Auto
country.dst Destination Country Auto
domain.src Source Domain Auto
Close
Select Meta Group


META GROUP DETAILS
GROUP NAME
Default Meta Keys
SHARING
Shared
DISPLAYED META KEYS
sessionid Session ID Hidden
size Data Size Hidden
payload Payload Size Hidden
medium Network Medium Hidden
eth.src Ethernet Source Address Auto
eth.dst Ethernet Destination Address Auto
eth.type Ethernet Protocol Auto
ip.proto IP Protocol Auto
ip.src Source IP Address Auto
ip.dst Destination IP Address Auto
ipv6.src Source IPv6 Address Auto
ipv6.dst Destination IPv6 Address Auto
tcp.srcport TCP Source Port Auto
Close
Select Meta Group

6. Do one of the following.
   a. To close the dialog, click **Close**.
   b. If you want to apply the meta group, click **Select Meta Group**.
      The dialog closes and the Filter Events panel is updated to reflect the meta keys in the selected meta group.

Select a Meta Group

1. With the Filter Events panel open in the Version 11.5 Events view, click the **Meta Groups** menu title.
   The menu drops down to display a list of meta groups and folders with a filtering option and a New Meta Group option. The list is sorted alphabetically and the name of the selected meta group is displayed in the menu label. This figure shows the menu after RSA HTTP was highlighted, but not selected.



2. Do one of the following:
   a. If the highlighted group is the one you want to apply, press **ENTER**.
   b. Begin typing text in the **Filter Meta Groups** field to search for a meta group name. As you type, the list is filtered to show only the meta group names that contain that string.
      When you see the group that you want to apply, click it or use the down or up arrow to highlight it, then press **ENTER**.
      The Filter Events panel is refreshed to include only meta keys in the selected meta group, and the menu title includes the selected group name. Your selection persists when you navigate away from the Events view.

Note: If a meta key in a meta group is not part of the selected service, it does not appear in the Filter Events panel or in the Events panel.

Create a Custom Meta Group

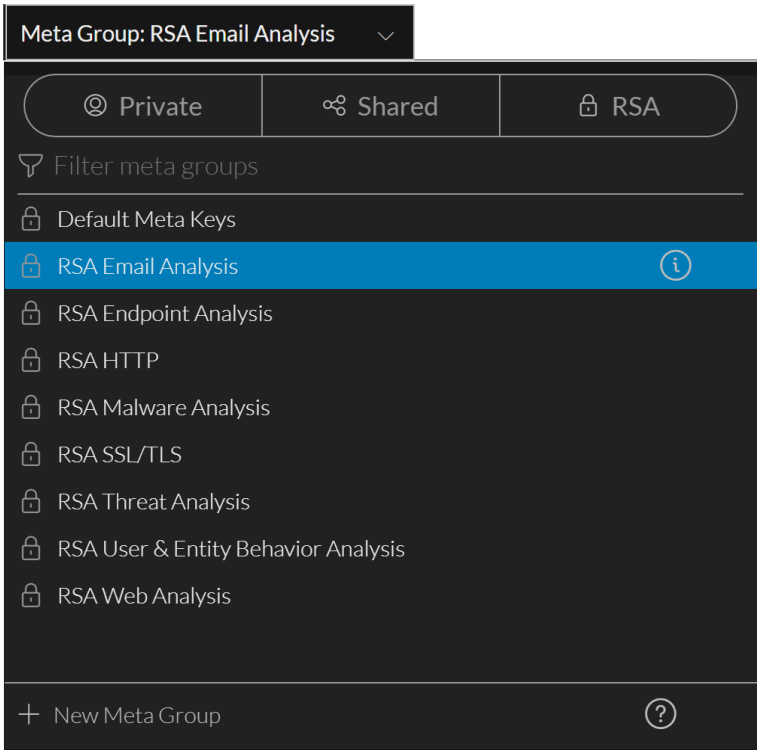
Custom meta groups must have a unique name up to 80 characters in length, and must have at least one meta key. If any other meta group has the name you type, whether shared or private, a message informs you that you need to use a different name. The Save Meta Group button is enabled when these criteria have been met. You can adjust the order of meta keys in a group by dragging keys in the Displayed Meta Keys list.

You can also set the initial view of each meta key: Open, Closed, Hidden, or Auto (the default setting).

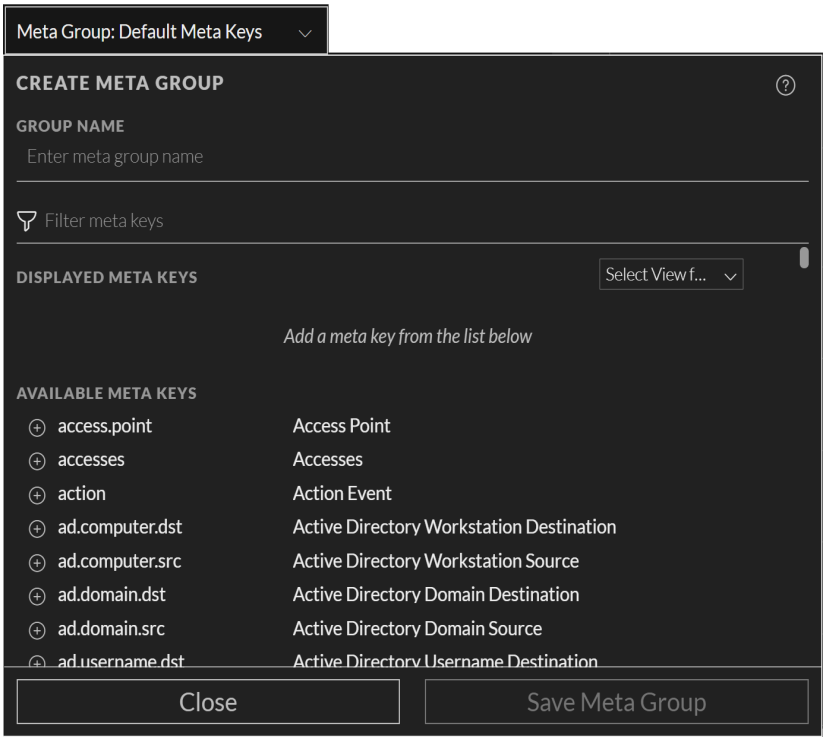
Note: You can also set the desired value for all meta keys at once. Make a note that changing the value of all meta keys might impact the performance.

- When set to Auto, the meta key is automatically loaded only if it is indexed, and non-indexed meta keys are Closed until opened manually. If you change the default view for a group of meta keys to Open and some of the meta keys are non-indexed, the non-indexed meta keys revert to Auto.
- Open meta keys are listed in the Filter Events panel, and the values are loaded.
- Closed meta keys are listed in the Filter Events panel, but the meta values are not loaded until you open the meta key.
- Hidden meta keys are not listed in the Filter Events panel at all. This is useful if you are using a single meta group for multiple purposes instead of creating several meta groups; you can turn off certain keys off without removing them from the meta group. You can also use the Hidden view when testing out some new keys or if you want to prepare a meta group with some new meta keys that are not yet available and would error out if in an Auto, Open, or Closed state.

1. With the Filter Events panel open in the 11.5 Events view, click the **Meta Groups** menu title.
   The menu drops down to display a list of meta groups and folders with the Filter Meta Groups field at the top and the + New Meta Group and Folder icon option at the bottom.

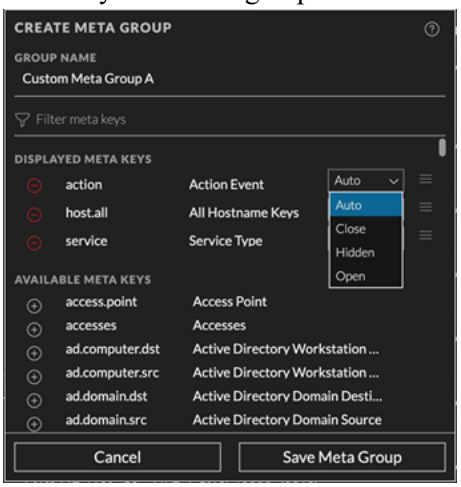


2. Select + **New Meta Group**.
   The Create Meta Group dialog is displayed.



3. In the **Group Name** field, type a unique name (maximum length of 80 characters) for the new meta group, for example, **Custom Meta Group A**.

4. If you want to share the new meta group with your organization, set the **Share with my organization** option.
5. To add a meta key to the meta group, select and add each meta key as follows:
   a. Type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Available Meta Keys** list.
   b. When you see the meta key that you want to add, click the add icon that precedes the meta key name.
   The meta key is added to the end of the Displayed Meta Keys list. (This list is also filtered using the text you typed.) The maximum number of meta keys in a meta group is 500. If you attempt to add another meta key when 500 are already included in the Displayed Meta Keys list, a message advises you that the group has the maximum number of meta keys.



6. (Optional) Next to each meta key, choose the initial view for the meta key: **Open**, **Close**, **Hidden**, or **Auto**.
7. (Optional) To find and remove a meta key from the meta group, type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Displayed Meta Keys** list. When you see the meta key that you want to remove, click the remove icon ( ) that precedes the meta key name in the **Displayed Meta Keys** list.
   The meta key is moved back to the Available Meta Keys list.
8. (Optional) To change the order of the displayed meta keys in the Displayed Meta Keys list, place the cursor over the list order icon ( ). When the cursor changes to the drag and drop icon ( ), drag the meta key up or down in the list.
9. Do one of the following:
   a. To close the dialog without creating the custom meta group, click **Cancel**.
   b. To create the group, click **Save Meta Group**.
   The new meta group is saved. If the new group is shared, it becomes available for all analysts. If it is private, only you can use the meta group. The buttons change to Done and Select Meta Group.

10. Do one of the following:
    a. To close the dialog, click **Done**.
    b. To close the dialog and select the new meta group, click **Select Meta Group**.
    The new group is added to the Meta Groups menu (in alphabetical order), and if you clicked Select Meta Group, the Filter Events panel is updated to show the meta keys and values in the new meta group.

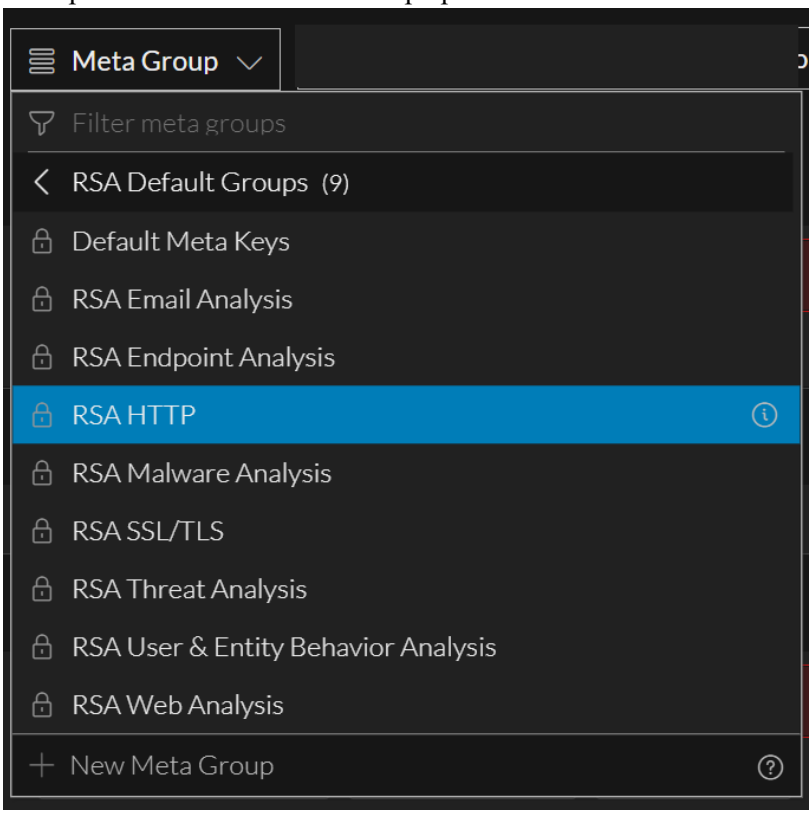
Delete a Custom Meta Group

You can delete any custom meta group, shared or private, that is not currently applied in the Events list and not used in a query profile. When you click the Delete button, a confirmation message allows you to confirm or cancel the deletion. If a meta group is being used in a query profile, the Delete button is disabled and a message identifies the query profile in which the meta group is used. The built-in meta groups are read only, and cannot be deleted.

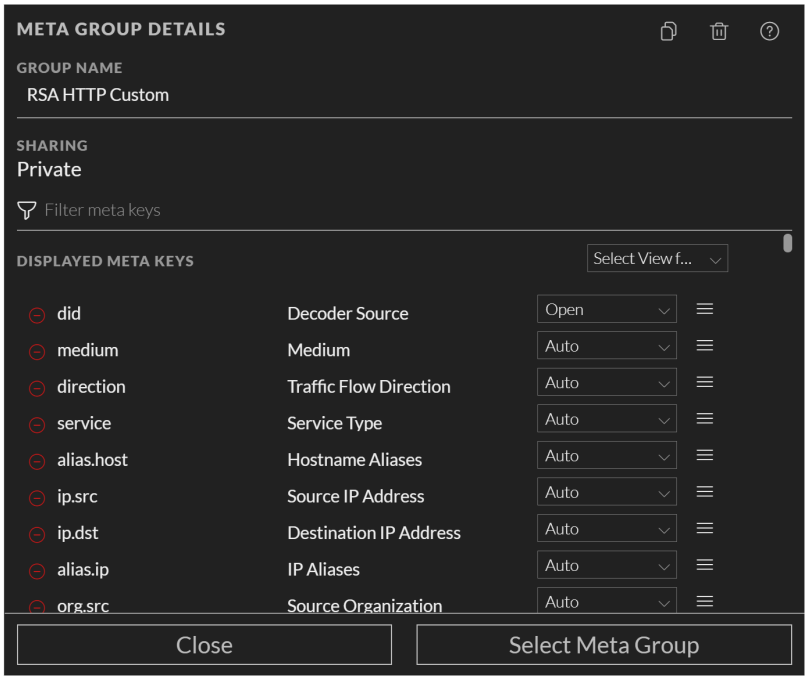
Caution: When you delete a shared meta group, the effect is global and the group is no longer available to any analyst.

To delete a custom meta group:

1. With the Filter Events panel open in the 11.6 Events view, click the **Meta Group** menu title.
   The menu drops down to display a list of meta groups and folders with the Filter Meta Groups field at the top and the + New Meta Group option at the bottom.



2. To delete a meta group, highlight a custom meta group and click the edit icon ( ) to the right of the name.
3. The Meta Group Details dialog opens with the details for the selected group displayed.

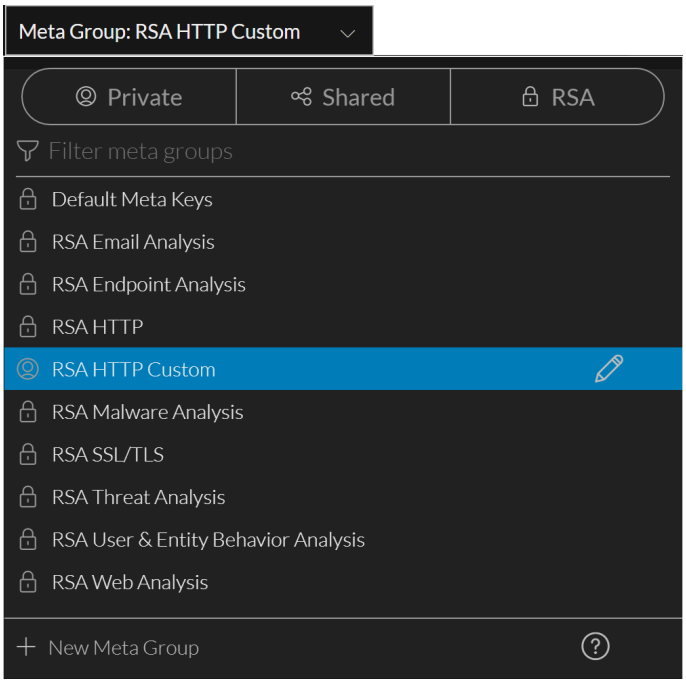


4. Click the delete group icon ( ).
   If the meta group is currently in effect, the following message is displayed: `This meta group cannot be deleted because it is currently active.`
   In Version 11.5, a confirmation message gives you the opportunity to confirm or cancel the deletion. Click **Cancel** or **Delete Meta Group**.
   The group is deleted and removed from the Meta Group menu. The meta group no longer appears anywhere for any analyst working in Investigate.

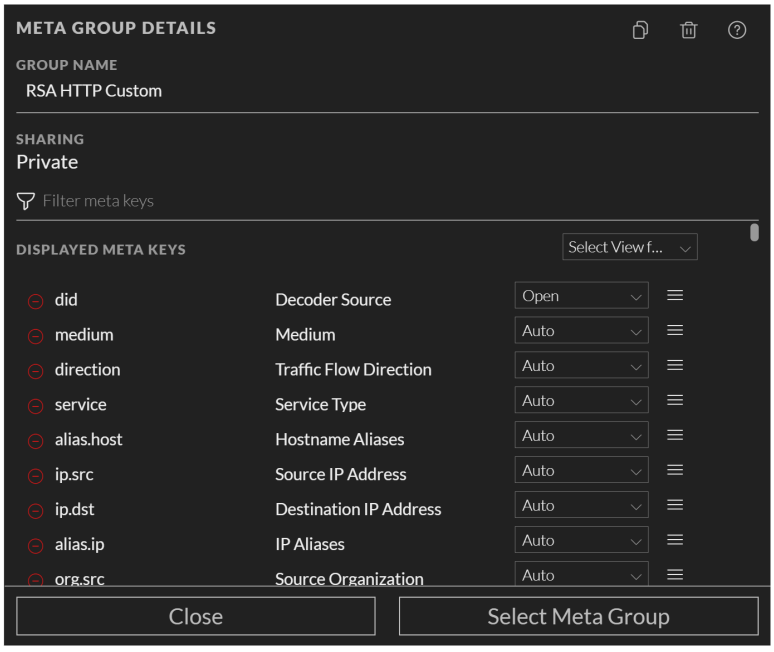
Edit a Custom Meta Group

You can edit a shared custom meta group, your own private meta group, a copy of a built-in meta group or a copy of live meta groups.

1. With the Filter Events panel open in the 11.5 Events view, click the **Meta Group** menu title and highlight the meta group that you want edit. This figure shows private column group RSA HTTP Custom highlighted with the edit icon is displayed to the right.

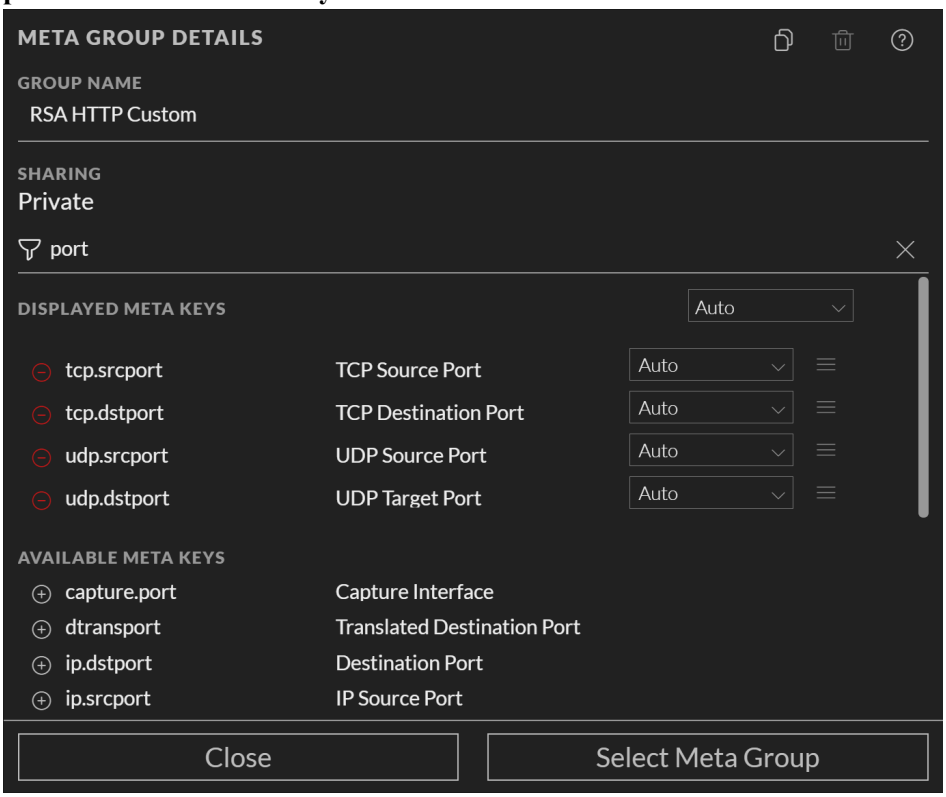


2. Click the edit icon ( ).
   The Meta Group Details dialog is displayed so that you can edit the location. You can add or delete meta keys and rearrange the order of the meta keys in the list.



3. (Optional) In the **Group Name** field, edit the name and location of the meta group.
4. (Optional) To add a meta key to the meta group, select and add each meta key as follows:
   a. Type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Available Meta Keys** list. Or just scroll through the list to find the meta key. For example, type

port in the **Filter meta keys** field.



b. When you see the meta key that you want to add, click the add icon that precedes the meta key name.

5. (Optional) To find and remove a meta key from the meta group, type a text string in the **Filter meta keys** field to look for meta keys that contain that text in the **Dispayed Meta Keys** list, or simply scroll through the list. When you see the meta key that you want to remove, click the remove icon ( ) that precedes the meta key name in the **Displayed Meta Keys** list.
   The meta key is moved back to the Available Meta Keys list.
6. (Optional) To change the order of the displayed meta keys in the Displayed Meta Keys list, place the cursor over the list order icon ( ). When the cursor changes to the drag and drop icon ( ), drag the meta key up or down in the list.
7. Do one of the following:
   a. To close the dialog without saving the changes to the custom meta group, click **Reset**.
   b. To save the edits to the meta group, click **Update Meta Group**.
   The updated meta group is saved, and the dialog is closed.

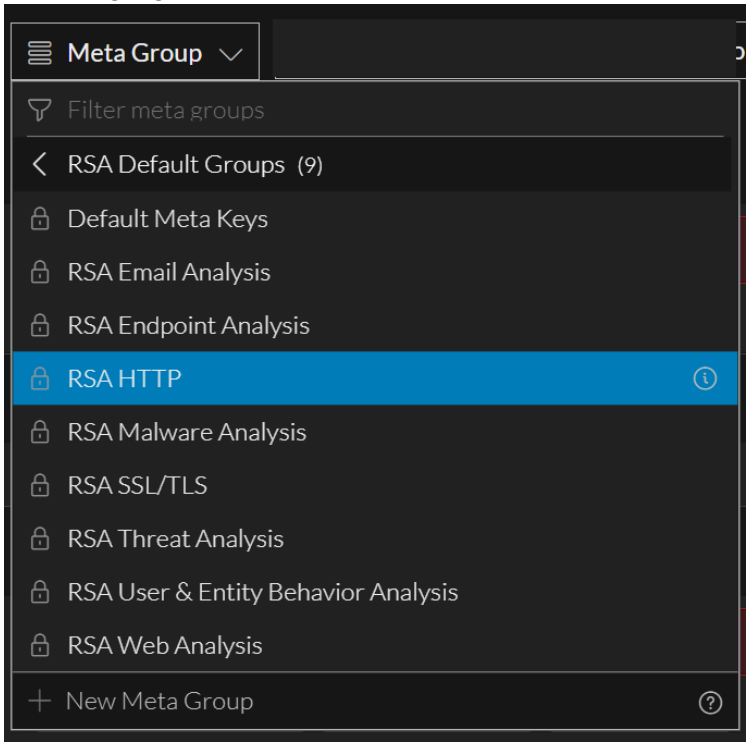
Copy a Meta Group (Version 11.5 and Later)

You can copy any meta group, built-in or custom, Live meta group, shared or private, as long as it does not have unsaved edits in progress. This is useful when you want a customized version of a built-in group. Also since you cannot change a custom group from private to shared or from shared to private, creating a copy allows you to select a different Sharing setting. When you copy a meta group, the same name is used with a number appended. For example, if you copy RSA HTTP twice, the first copy is named RSA HTTP-1, and a second copy is named RSA HTTP-2. After you copy the group, you can edit the copy to give it a new name and manage meta keys in the group.

Note: Some meta groups created in the Legacy Events view may have more 500 meta keys, which is above the limit for meta groups in the Events view. If you copy a group with more than 500 meta keys, you must remove the excess meta keys when you edit the meta group.

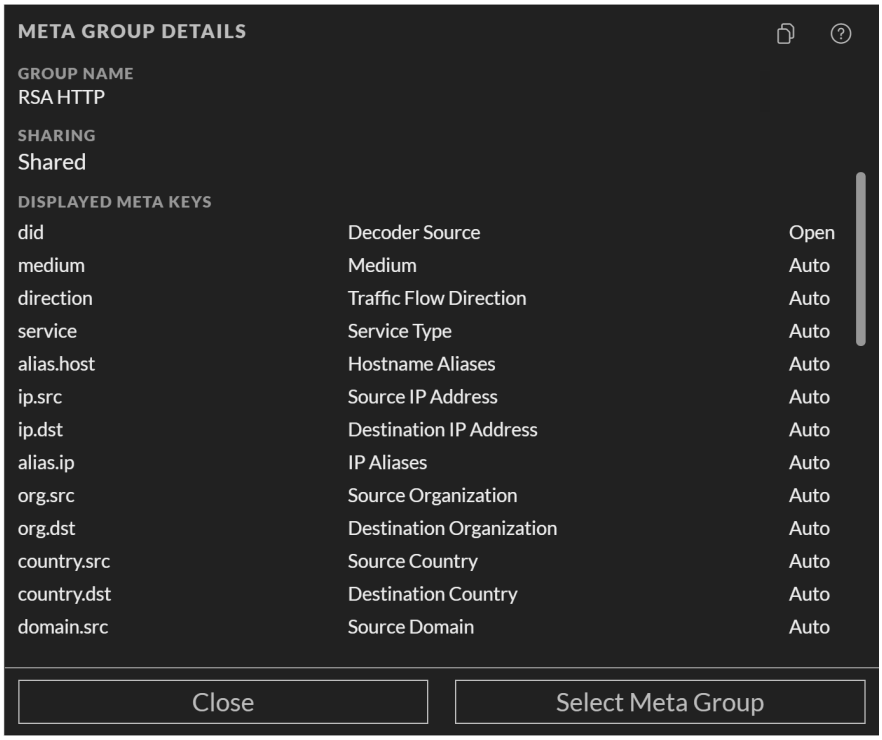
To copy a meta group:

1. With the Filter Events panel open in the 11.6 Events view, click the **Meta Group** menu title.
   The menu drops down to display a list of meta groups and folders.
2. Highlight the meta group that you want copy.
   If you highlighted a built-in meta group, the information icon ( ) is displayed to the right. If you highlighted a custom meta group, the edit icon ( )is displayed to the right. This figure shows RSA HTTP highlighted.

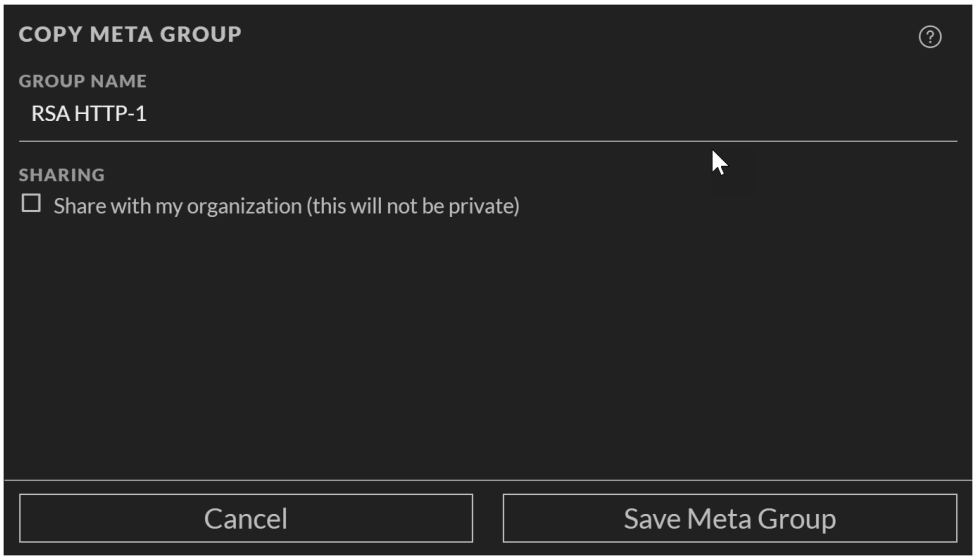


3. Do one of the following:
   a. Click the information icon ( ).

b. Click the edit icon ( ).
The Meta Group Details dialog is displayed. This figure shows the dialog for a built-in group.

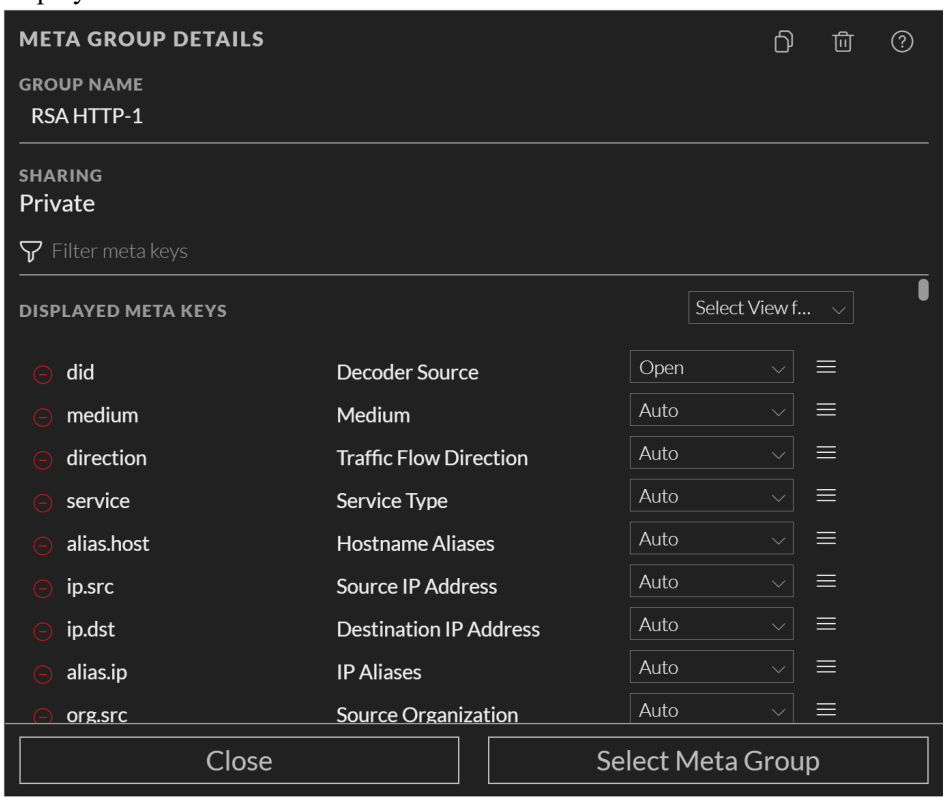


4. Click the Copy icon ( ).
The Copy Meta Group dialog is displayed with a `-n` appended to the original meta group name.



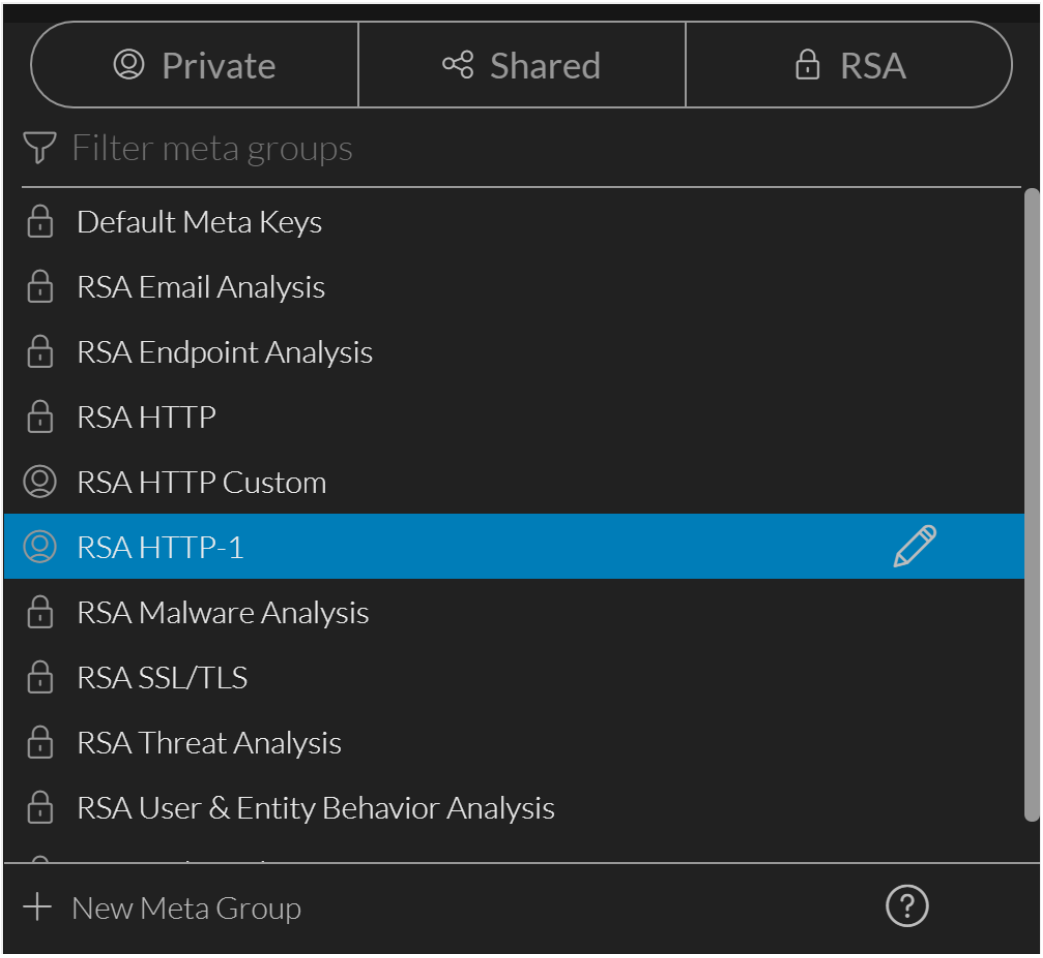
5. (Optional) In the **Group Name** field, edit the name and location of the meta group.
6. Do one of the following:
   a. To close the dialog without copying the group, click **Cancel**.
   b. To save the copy of the meta group, click **Save Meta Group**.
   The copy of the meta group is saved, and the Meta Group Details dialog for the copied group is

displayed.



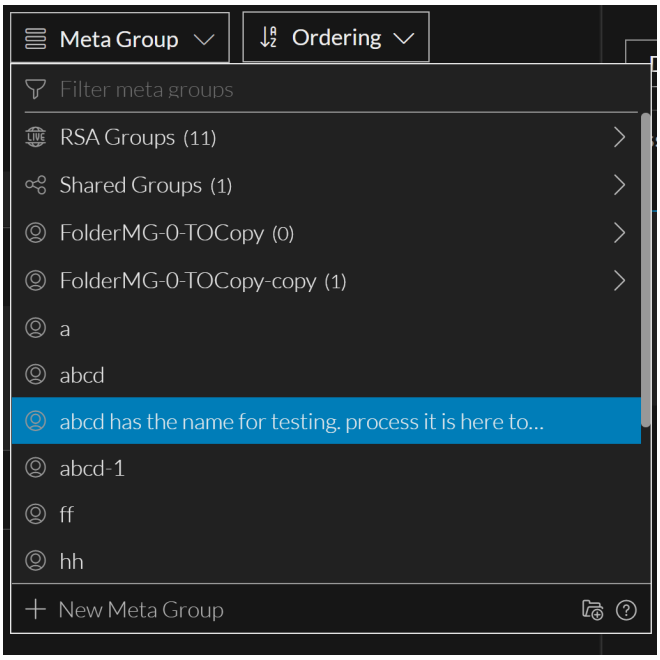
7. Do one of the following:
   a. To close the dialog without editing, click **Close**.
   b. To close the dialog and select the copy of the meta group, click **Select Meta Group**.
   The group is added to the Meta Group menu. The figure below has a private copy of the RSA

HTTP meta group.



Meta Group Folders

Users can create editable Shared and Private group folders. The contents of a private group folder and their contents are displayed outside RSA Groups and Shared group folders. For example, the below image shows private content below the Shared Groups folder.

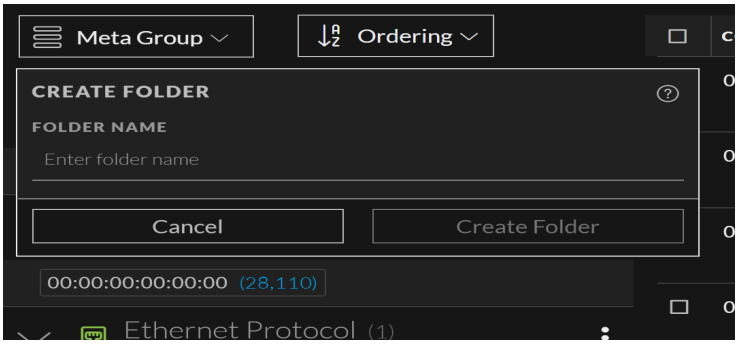


This section describes how to add, edit, import, export, copy, and delete custom meta groups and folders.

Create a Meta Group Folder

You can create meta group folders as shared and private folders. And, if the folder name already exists then you are prompted to provide a unique name.

1. With the Filter Events panel open in the Events view, click the Meta Groups menu title. The menu drops down to display a list of meta groups and folders.
2. Click  .
   The Create Folder dialog is displayed.

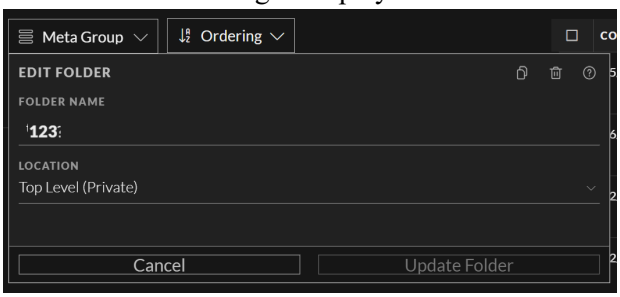


3. In the **Folder Name** field, type a unique name for the new meta group folder.
4. Click **Create Folder**.

Edit and Move Meta Group Folder

After you create a meta group folder you can edit or move it, however the folders inside RSA Groups (RSA Live content and RSA OOTB Groups) cannot be edited and moved. The folders inside private and shared folders can be edited and moved only within their respective groups. For example, you cannot move a shared folder into a private folder and vice-versa.

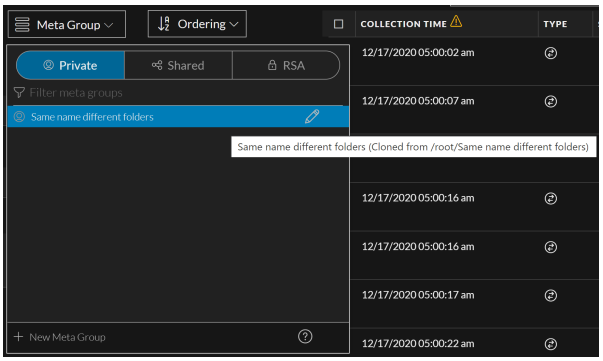
1. With the Filter Events panel open in the 11.6 Events view, click the Meta Group menu title and highlight the meta group that you want edit.
2. Click .
   The Edit Folder dialog is displayed.



3. In the **Folder Name** field, type a unique name for the meta group folder.
4. Select the location of the folder to be edited.
5. Click **Update Folder**.

Copy Meta Group Folder

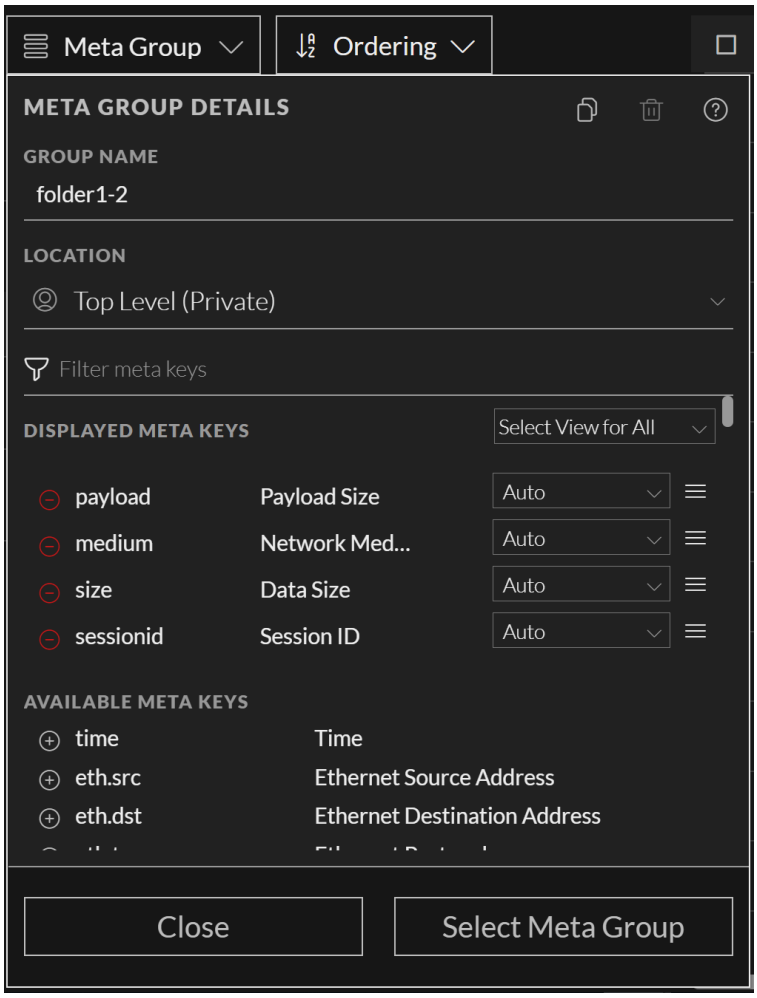
Users can copy any type of meta group folders namely - RSA , Shared and Private. However, by default for RSA groups copy folder will create a copy in the private section (root level) but can change the location of the folder to a shared folder or any other private folder. The meta group can be copied by clicking the clone icon ( ). After copying, the meta group folders are displayed selected location (Shared or Private category). You can hover over on the copied item to view a tooltip that indicates the path from which the meta group has been copied. In case you need to search for a specific meta group, you can type the name of the meta group in the filter field ( ) at the folder level and the meta group will be filtered from the selected folder.



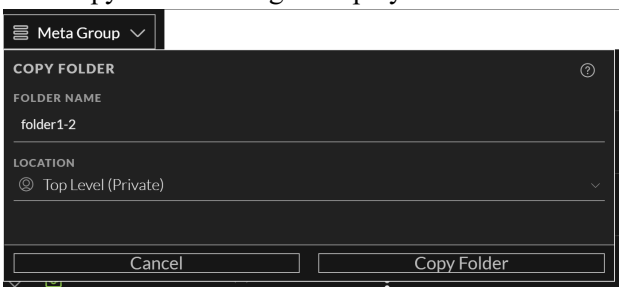
Copying Meta Group Private or Shared Folders

You can copy meta group folders from RSA groups to private and RSA groups to shared, private to shared, private to private, shared to shared and shared to private groups. When you copy a folder the content inside it gets copied except the sub-folders. When you copy a private folder into a shared folder, the folder and its content no longer remain private.

1. With the Filter Events panel open in the Events view, click the Meta Group menu title. The menu drops down to display a list of meta groups and folders.
2. Select a folder you want to copy.
3. Click edit and then click copy .



The Copy Folder dialog is displayed.



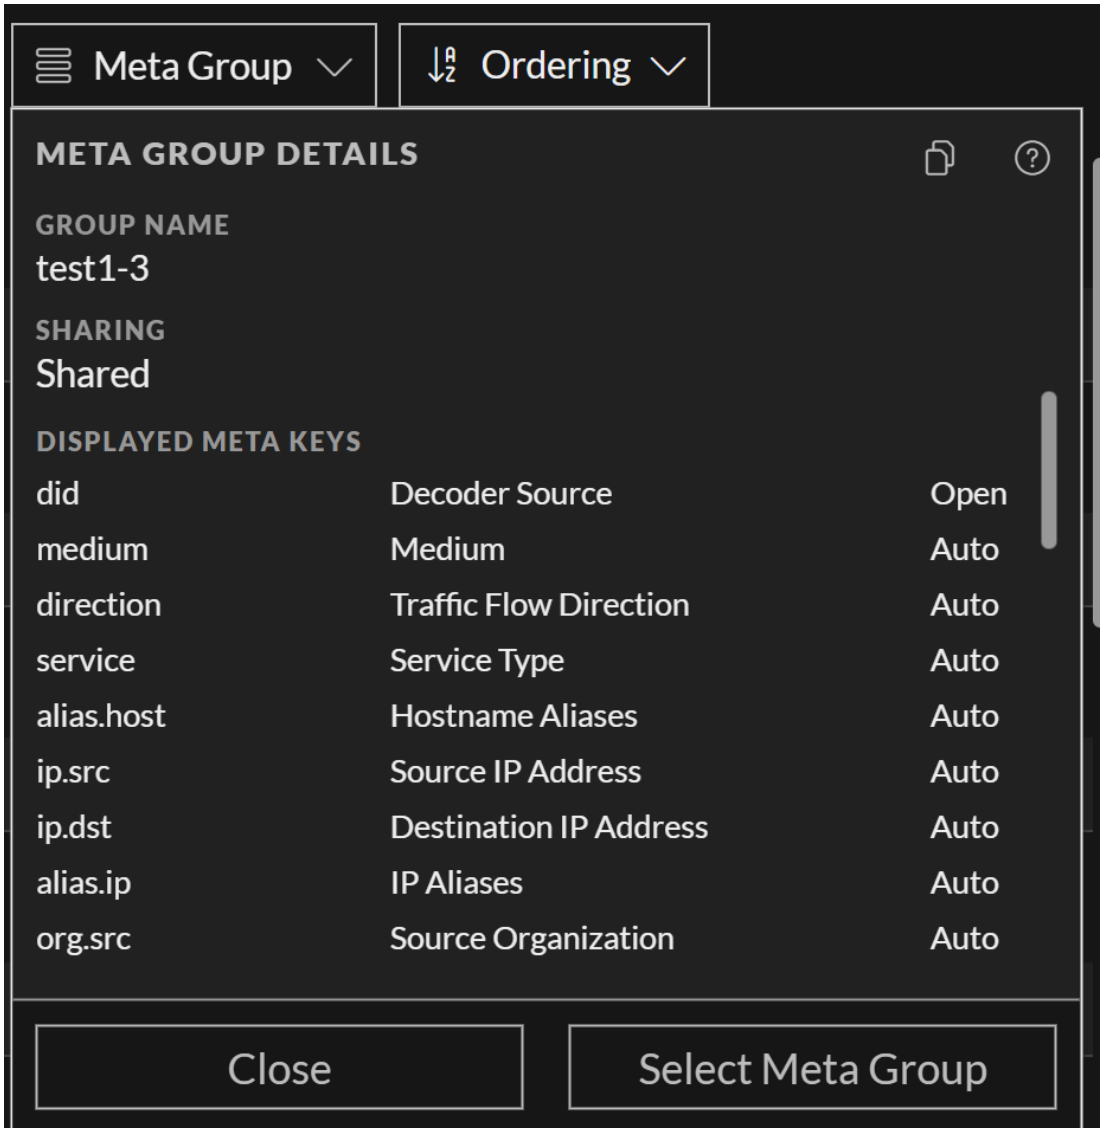
4. In the **Folder Name** field, type a unique name for the new meta group folder.
5. Select the location of the folder to be copied.
6. Click **Copy Folder**.

Copy Meta Group Folder Deployed from Live

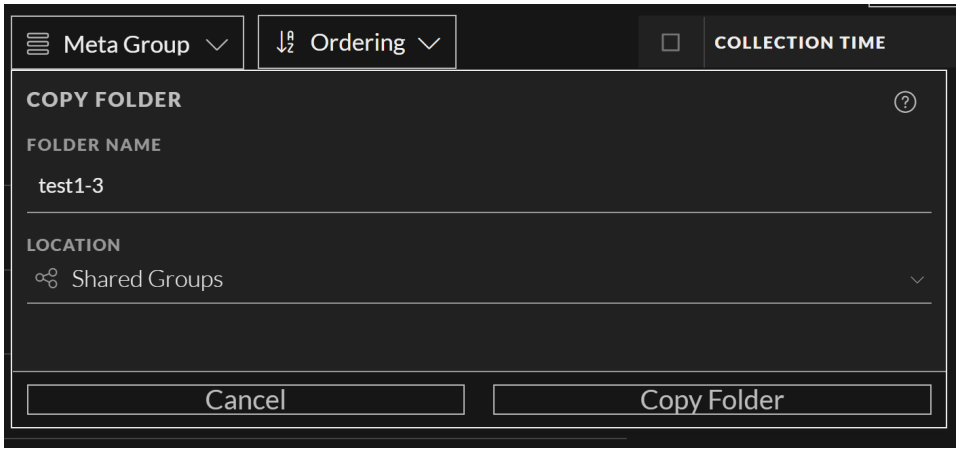
You can copy meta group folder deployed from Live located under RSA Groups category to any other location like Shared groups or to a private folder.

1. With the Filter Events panel open in the Events view, click Meta Group menu title. The menu drops down to display a list of meta groups and folders.
2. Click on RSA groups and select a Live Meta Group folder you want to copy.

3. Click  and the click copy  .



The Copy Folder dialog is displayed.

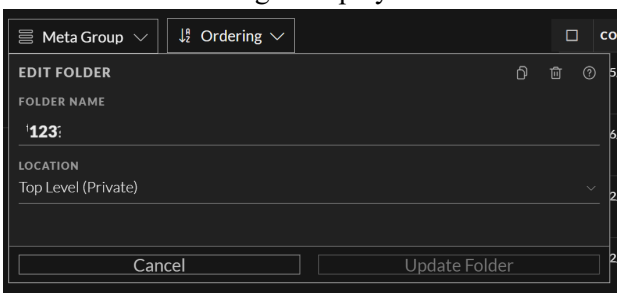


4. Select the location of the folder to be copied.
5. Click **Copy Folder**.
   The folder and the first level contents are copied, sub-Folders are not copied. The copied Meta Group Folder and its contents are displayed as the original meta group name appended with a -n .

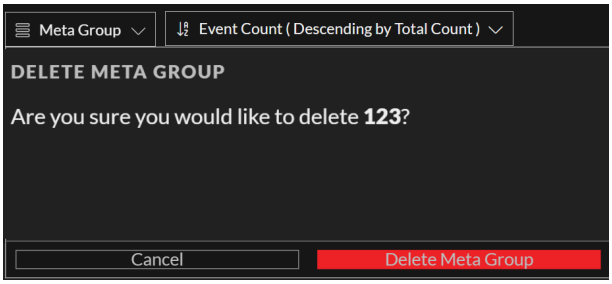
Delete Meta Group Folder

If you don't want to retain a folder you can delete it. However, once the folder is deleted it cannot be retrieved.

1. With the Filter Events panel open in the Events view, click the Meta Group menu title. The menu drops down to display a list of meta groups and folders.
2. Select a folder to be deleted.
3. Click edit .
   The Edit Folder dialog is displayed.



4. Click delete .
   A warning message is displayed to confirm the action.

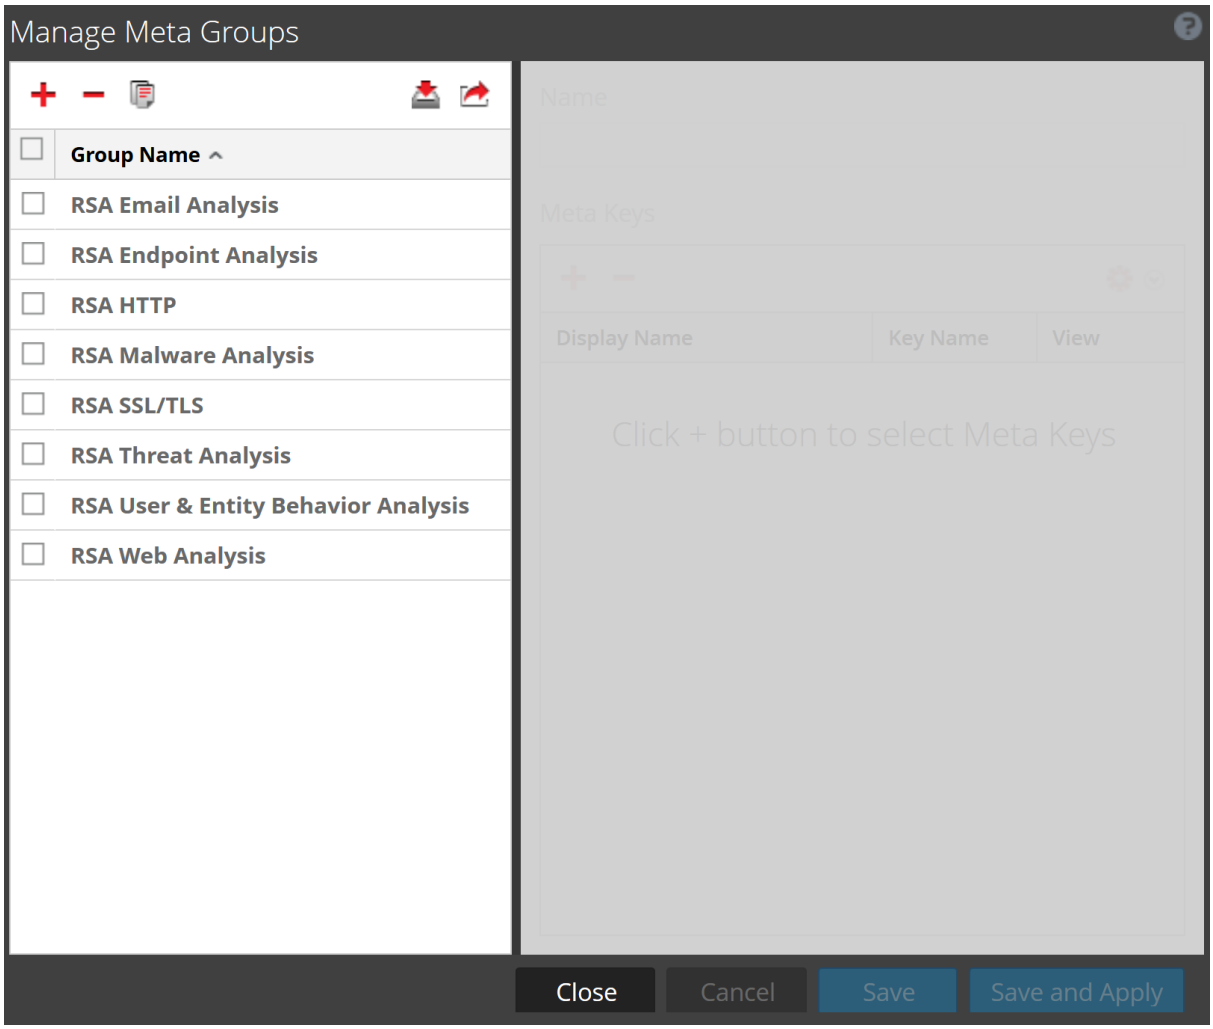


5. (Optional) Select the checkbox, if you want to delete the folder along with all the contents inside the selected folder.
   If you do not select the checkbox, then the content will be moved to the parent folder after the required folder is deleted.
6. Click **OK** to delete.

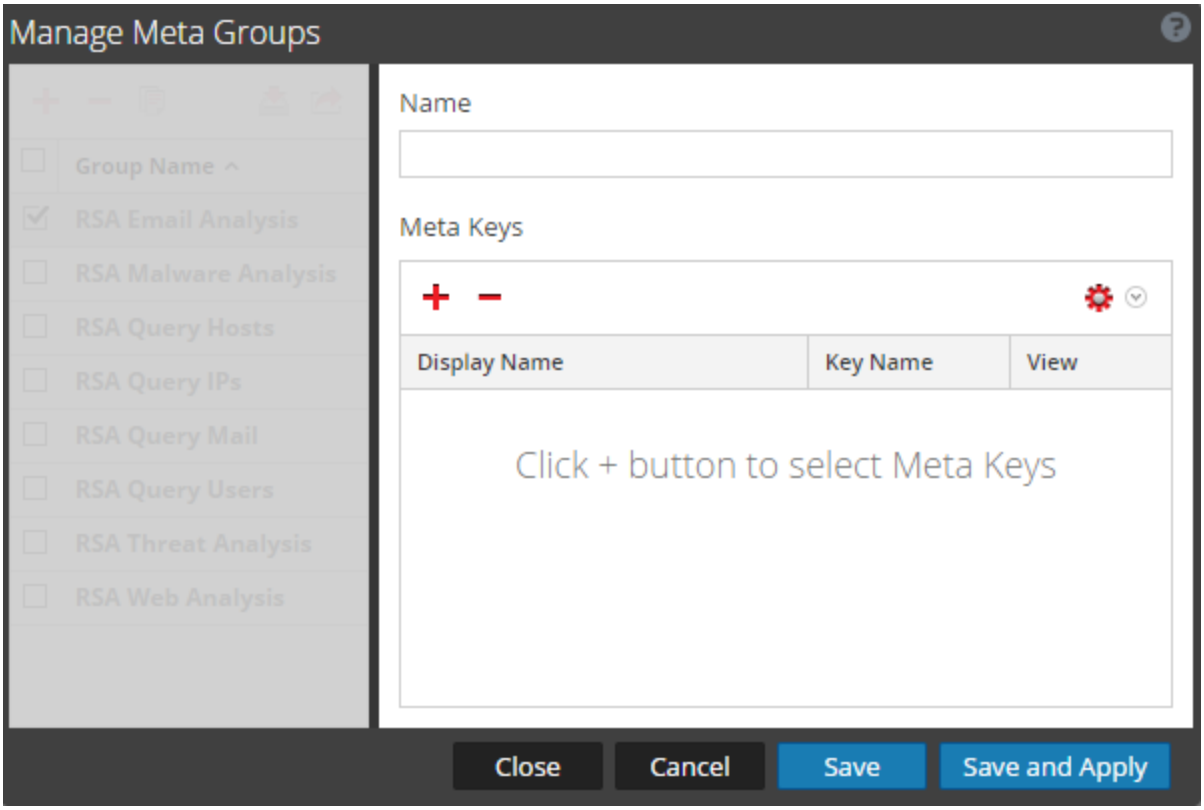
Work with Meta Groups in the Navigate View

Create a Meta Group and Add Meta Keys

1. While investigating a service in the **Navigate** view, select **Meta > Manage Meta Groups** in the toolbar.
   The Manage Meta Groups dialog is displayed. Initially only built-in groups are configured for a service and listed under Group Name. If other custom groups have been configured, they are also listed under Group Name.

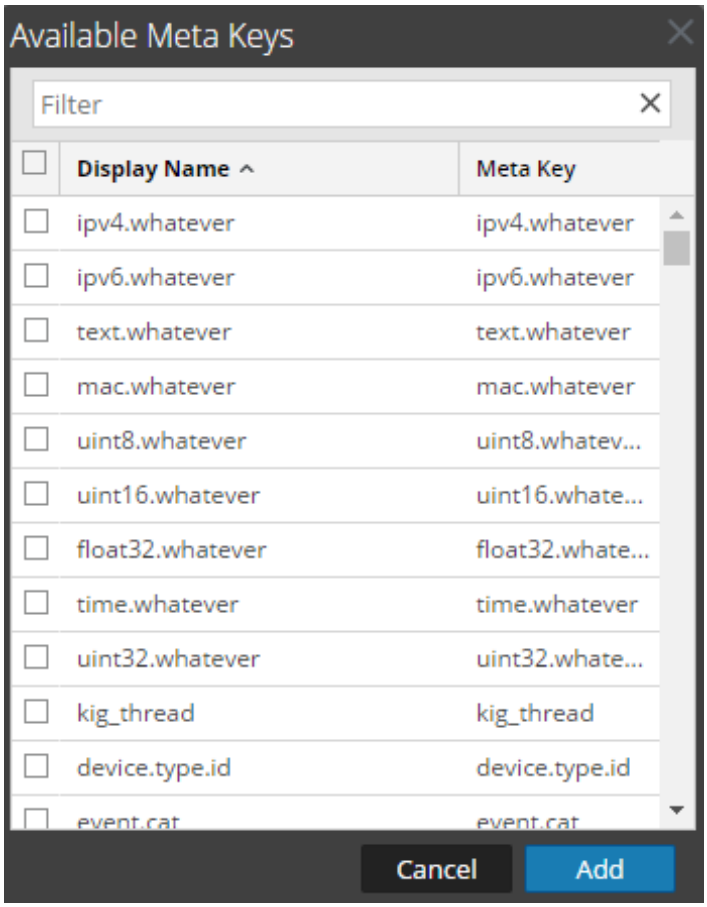


2. In the toolbar at the top of the Meta Groups list, click +.
   The form to the right opens for editing.



3. Type a name for the new meta group in the **Name** field.

4. In the **Meta Keys** toolbar, click +.
   The Available Meta Keys dialog is displayed, with keys in alphabetical order.

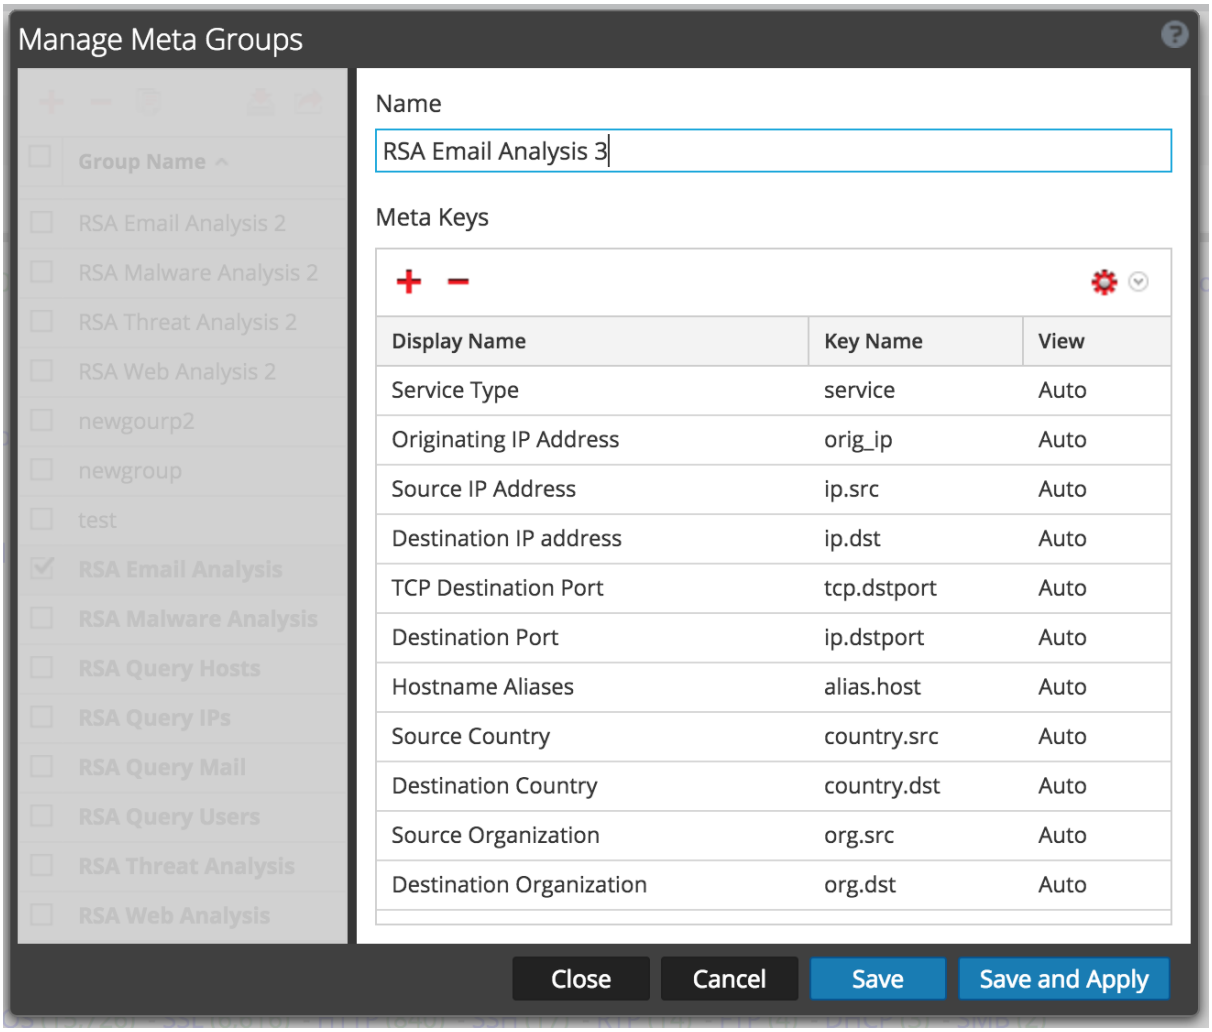


5. To filter the list of meta keys, type a word or phrase in the **Filter** field and press **Enter**.
   The list displays matching meta keys based on a case-insensitive search. Delete the filter text and press **Enter** to remove the filter.
6. To select individual meta keys to include in the meta group, select the checkboxes. To select all meta keys, select the checkbox in the title bar and click **Add**.
   The selected meta keys are added to the meta keys list.
7. (Optional) If you want to change the order in which the meta keys load and are listed in an investigation, click and drag one or more meta keys to a new position.
8. To finish creating the meta group do one of the following:
   a. To save the meta group, click **Save**.
      The group is created and available for use.
   b. To save and apply the meta group to the current Investigation view, click **Save and Apply**.
      The group is created and applied immediately to the current Investigation view.
9. Click **Close**.

Copy and Edit a Meta Group

If you want to customize a built-in meta group, you need to duplicate the group and then edit the duplicate.

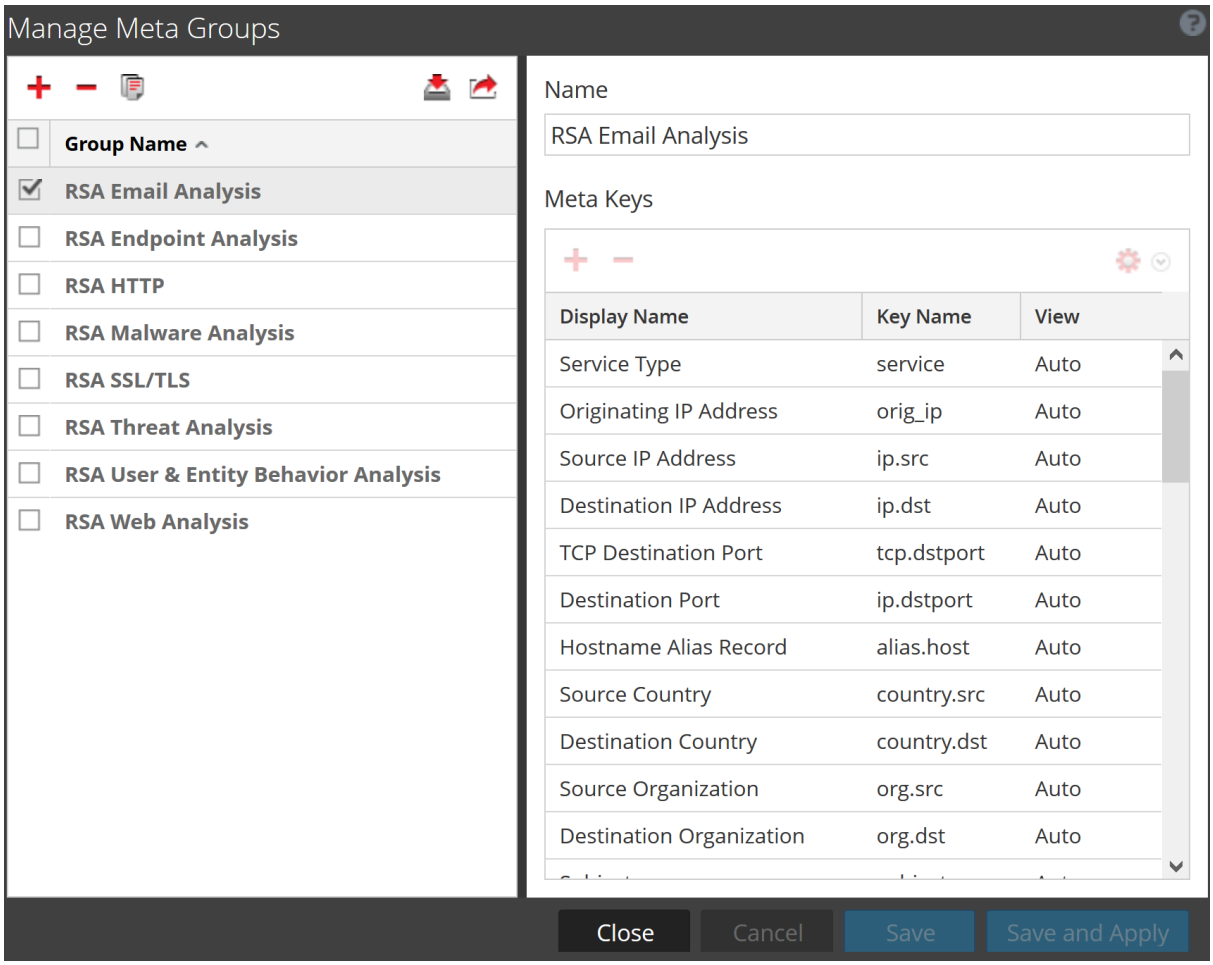
1. Select a built-in meta group from the Manage Meta Groups list and click .
   The form to the right opens for editing with all of the meta keys as they are in the built-in group.



2. Enter a name for the new group and continue editing as described in "Edit a Meta Group" below.

Edit a Custom Meta Group

1. Select a custom group from the **Meta Groups** list.
   The form to the right opens for editing.



2. (Optional) Edit the Name of the group.
3. (Optional) Add new meta keys, as described above in "Create a Meta Group and Add Meta Keys."
4. (Optional) To set the order for the keys, drag and drop one or more keys.
5. (Optional) To change the initial view of a meta key, click  and choose one of the possible views.
   When you modify the meta group, you cannot set the key to OPEN. If you change the default view for a group of meta keys to OPEN and some of the meta keys are non-indexed, the non-indexed meta keys revert to AUTO. As a result, the meta key is automatically loaded only if it is indexed, and non-indexed meta keys are CLOSED until opened manually.
   The value for the initial view is displayed in the View column.
6. To save, the changes, click **Save**.
7. To apply the changes to the current Navigate view, click **Save and Apply**.

Delete a Meta Group

1. In the **Meta Groups** list, select the group to be removed.
2. Click −.
   A confirmation dialog provides an opportunity to cancel or complete the request.
3. Click **Yes**.
   The meta group is deleted. When you close the window, if the deleted group was the currently applied meta group, it is removed and the default meta keys are used to build the view.

Export a Meta Group

User-defined meta groups are created on individual services. To make meta groups available to another service, you must export them to your local file system. To export one or more meta groups:

1. In the **Meta Groups** list, select one or more groups to be exported.
2. Click  .
   The selected groups are downloaded to your local file system as a **MetaGroups.jsn file**. Every download of meta groups has the same name with a numeral appended to avoid overwriting previous downloads.

Import a Meta Group

To make user-defined meta groups from another service available to the currently investigated service, you must import the `MetaGroups.jsn` file from the local file system. When you import meta groups, an error message is displayed if any of the groups are already present. To import a group that is a duplicate, you must first delete the existing group. If you want to delete a meta group, it cannot be in use by a profile.

To import meta groups:

1. In the **Meta Groups** list, select a file to import and click  .
   The selection dialog is displayed.

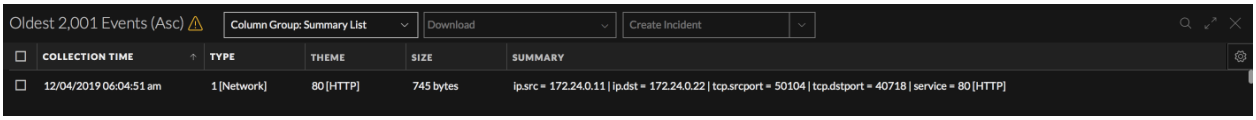


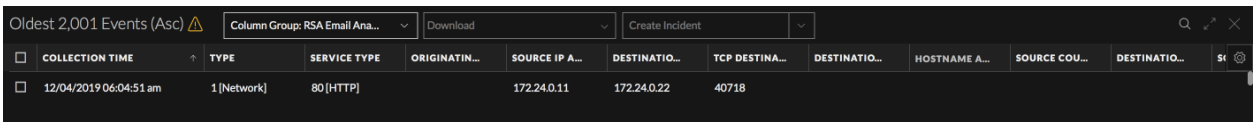
2. Click **Browse** and navigate to the directory on your local file system where the downloaded `MetaGroups.jsn` files are stored. Select a file and click **Open**.
   The filename is displayed in the Upload File field.

3. Click **Upload**.
   The upload process begins, and a message indicates that the upload was successful. The meta groups are added to Meta Group list. If the file is a duplicate of an existing meta group, a dialog tells you that the meta group already exists.

Use Columns and Column Groups in the Events List

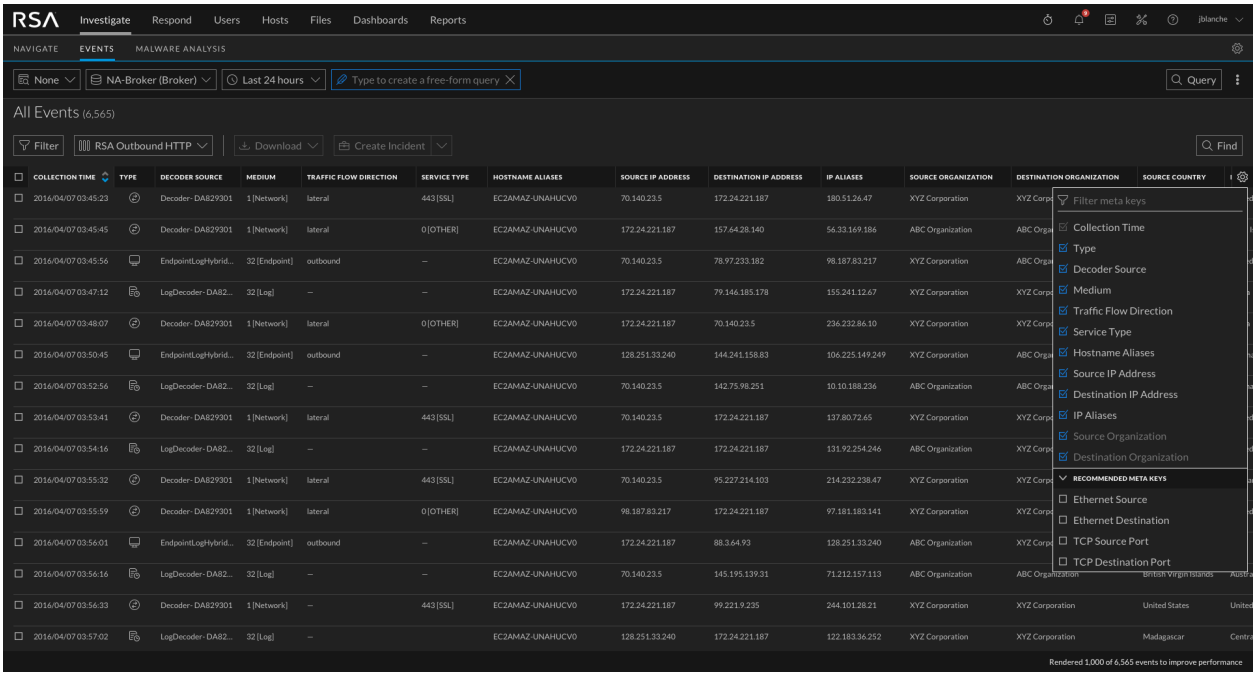
When the events list in Investigate is populated with events, each column lists the values returned for a meta key. Changing the meta keys displayed in the events list is a useful method of narrowing the focus of your investigation. For example, compare these two figures showing the same set of events with different columns. The first figure has five columns, **Collection Time**, **Type**, **Theme**, **Size**, and **Summary**. These are just the basic information, not specialized in any way. The second figure has many more columns that contain information useful when investigating email; you can scroll to the right to see the additional columns.





You can adjust the events list as you work, selecting different columns to be displayed, rearranging the order of the columns, changing the width of the columns, and choosing a column by which the list is sorted. Manual adjustments are easy to make if you know which meta keys are relevant. Manual adjustments apply only to the current session in Version 11.5; in Version 11.5.1, the column width is an exception. When you adjust the column width, it is preserved as a personal preference and is applied every time the column is used in the Events list, overriding any default column width.

In version 11.6, you can select additional columns with data for the meta keys that you are viewing. This will enable you to obtain all the relevant meta key information from the Filter Events panel however the recommendations can change based on the selected meta groups. The following figure displays the additional meta key information under the Recommended Meta Keys section.



To improve your ability to see relevant meta keys quickly when looking at events in the Legacy Events view and the Events view, you can change the set of meta keys displayed by applying a column group. A column group defines the meta keys or meta entities that are displayed as columns, the position of the column in the Events list, and the default width of the column. A column group must have at least one column. Column groups are useful in themselves, and they become even more useful when you combine them with meta groups and preQueries to define query profiles (see Use Query Profiles to Encapsulate Common Areas for Investigation).

The same column groups are shared between the Legacy Events view and the Events view. When importing a column group, the imported group is limited to the available meta keys for the service being investigated. Private column groups created in the Events view are not available in the Legacy Events view or for use in Query Profiles in the Navigate view.

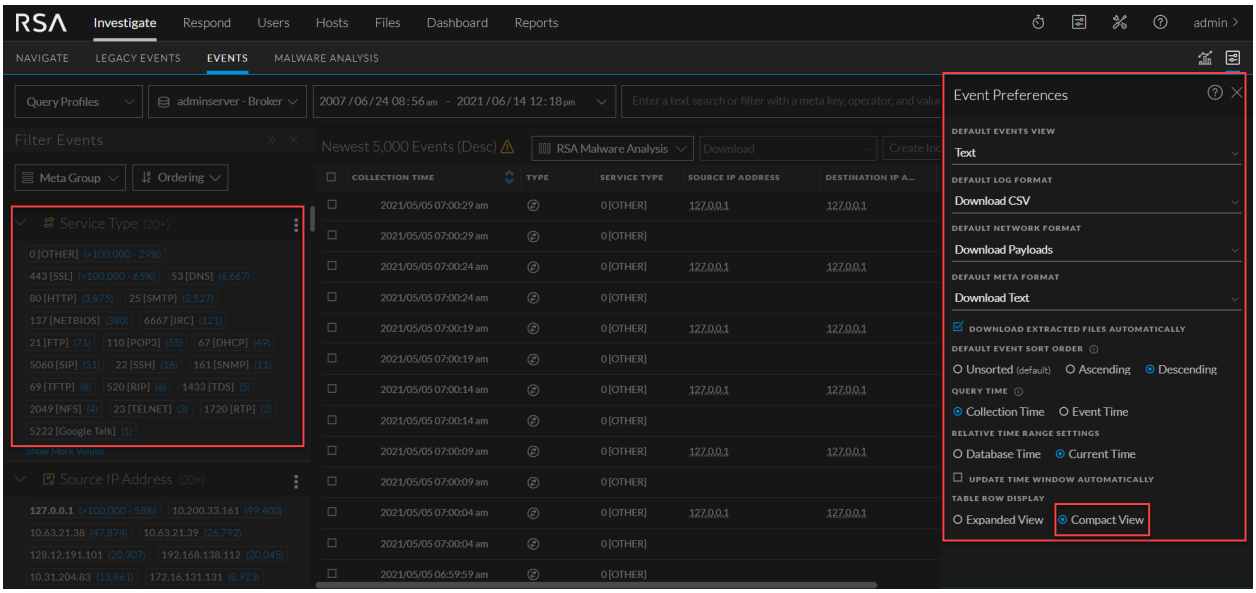
Note: In the Navigate view and Legacy Events view, you can manually add non-indexed meta keys (or keys that are not in the index at all) to a meta group or column group. The non-indexed meta keys are fully available (manageable and displayable) in the Navigate view and Legacy Events view, but only partially (displayable in the Filter Events panel) in the Events view. The Events view Filter Events panel can display data for non-indexed meta keys that are already included in a meta group, but you cannot add non-indexed meta keys while you are editing a meta group. The non-indexed meta keys in a column group do not display data in a column and new non-indexed meta keys cannot be added to a column group in Events view.

Large column groups can have a performance impact when loading data because the values for each meta key are loaded in the events list. To minimize impact on performance, the Events view has a fixed limit on the number of meta keys in a column group. The maximum number of meta keys in a column group is 40. (Because several default meta keys are included you may see a few more than 40 displayed on the screen.) Meta keys that are not in the selected column group are not loaded in the events list. By default we load all columns in the group, but only 15 are displayed by default.

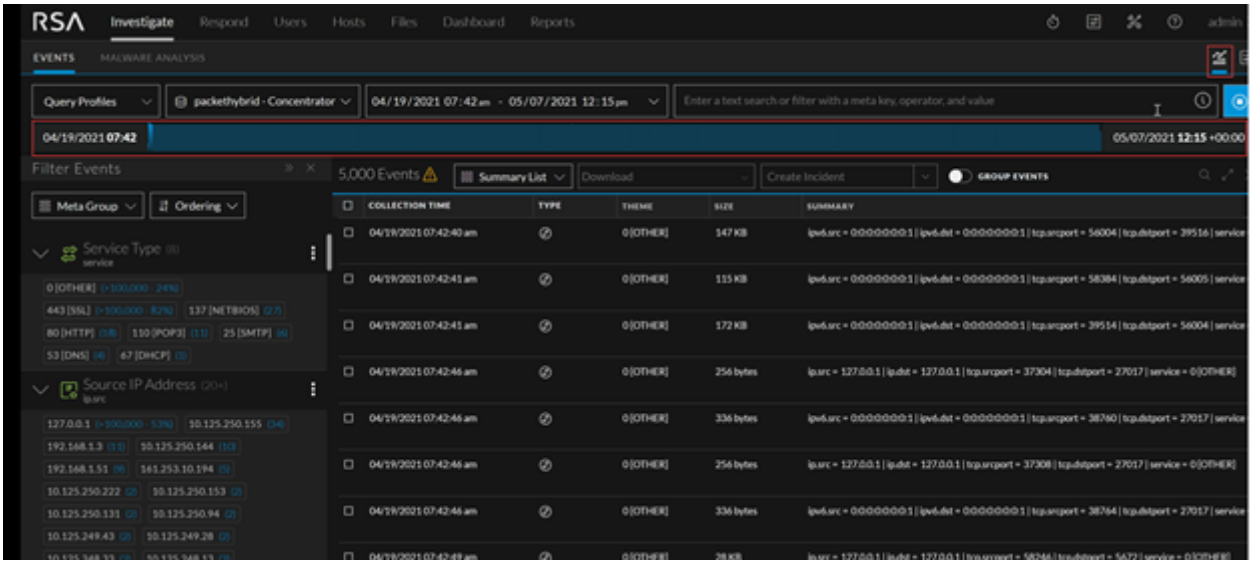
The Legacy Events view does not have a limit on the number of meta keys in a column group, and may have more than 40 meta keys in a column group. If you apply a column group with more than 40 meta keys that was created in the Legacy Events view, all columns are loaded in the Events view. If you copy a group with more than 40 columns, you must remove the excess columns when you edit the column group.

Note: In Version 11.3, column groups are created and managed in the Events view and available for use in the Event Analysis view. All existing column groups, both built-in and custom, are available in the 11.4 Events view. The complete column group management functionality is available in the Legacy Events view, and all functionality except cloning, importing, and exporting column groups is available in the 11.4 Events view. In Version 11.5, cloning is also available in the Events view, but importing and exporting are not.

In 11.6.1, the Investigate > Events Preferences view has been added to make optimum use of the space to enable analysts to view maximum details related to the events they are analyzing.

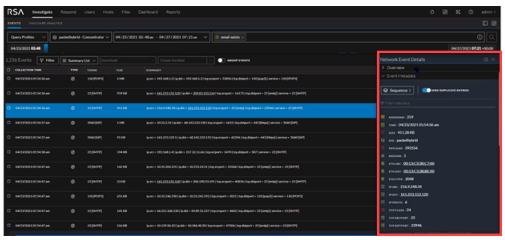


Additionally, an analyst can now view the timeline details of an event by clicking the ▣ icon. When clicked the time line that displays the date and time range s displayed as shown in the image.

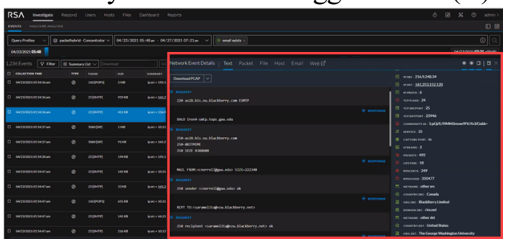


The Investigate > Events reconstruction panel has been modified to display an overlay that will contain an Overview tab and a Meta panel tab that can be expanded or collapsed. This will enable the analyst to view the headers and meta panel of the events optimally. The analysts can also toggle **Hide Duplicate Events** option and view only the relevant details of a selected event.

When an analyst navigates to this page the following view will be displayed.

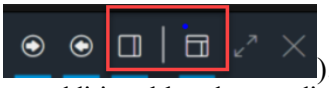


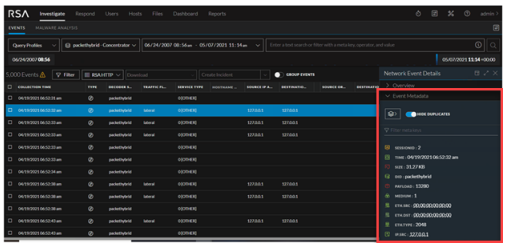
The analyst can use the toggle button ( ) to view the details related to the selected event.



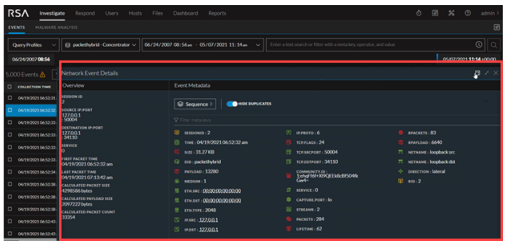
The binocular icon is enabled only if an event payload is open for a search related to a selected event.

The Overview Tab displays all the headers related to a specific event and the Event Metadata Panel displays all the metadata related to the selected event.

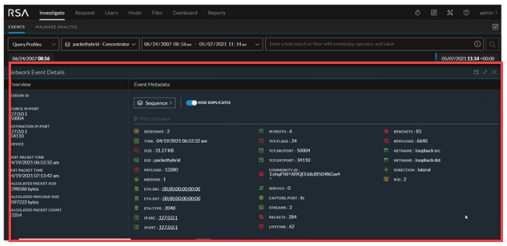
In 11.7, when you open the meta panel to view the selected event details, all the available headers will be displayed. The meta panel contains Expand buttons (). When you click on an event, the meta panel is displayed. In case, there are no additional headers to display, an error related to the header error is displayed.



The analyst can use the expand option to broaden the meta panels in three different views. When the outward facing arrow is clicked, it expands all the details displayed under the Overview and Event Metadata panel are displayed.



When clicked one more time, you can get an expanded view.



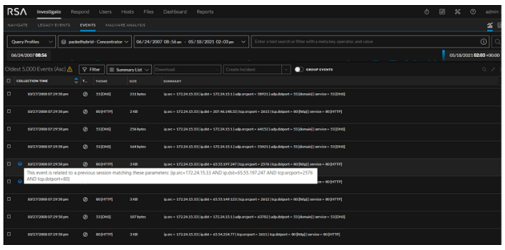
The analyst can revert the screen display by clicking the icon with the inward facing arrow , so that the window reverts to the earlier position. This helps the analyst to view the details of each event in an optimum manner.

The analysts can search for related sessions for a specific event as part of an investigation. The search for related sessions can be performed by navigating to the **Investigate** > **Events** page. You can click the icon and select either Find Related Sessions or Find Related Sessions in New Tab option from the drop-down.

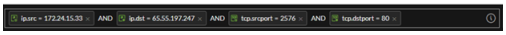


When you select the **Find Related Sessions** option, all the events that are matched by selected event's query will be displayed in the current window. And if you select the **Find Related Sessions in New Tab** option the results are displayed in a new tab. The analyst can further investigate on each of the related session.

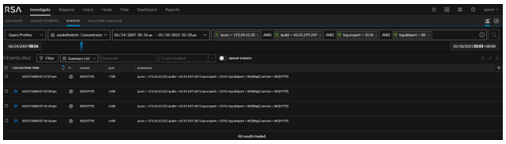
The query is based on the information displayed in the hover over text of an event. For example, in the below image the event has two split sessions where one event is split into one another session.



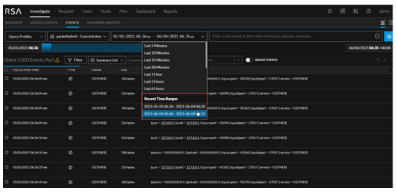
So, in this case when a search is done on these parameters, the related session for the following query is displayed.



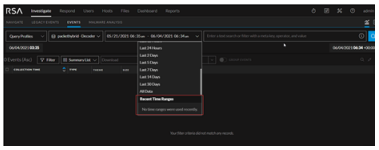
The results are displayed in the following format:



An analyst can view the last five time-ranges that were used recently, as the selection will be saved and displayed under the Recent Time Ranges section. The selection is saved separately for a user per service. For example, for Concentrator service, if you select Last 30 Days as the time-range, this will create an entry under the Recent Time Ranges section for Concentrator and when you select the same service in the next session, then the time range will be shown as below:

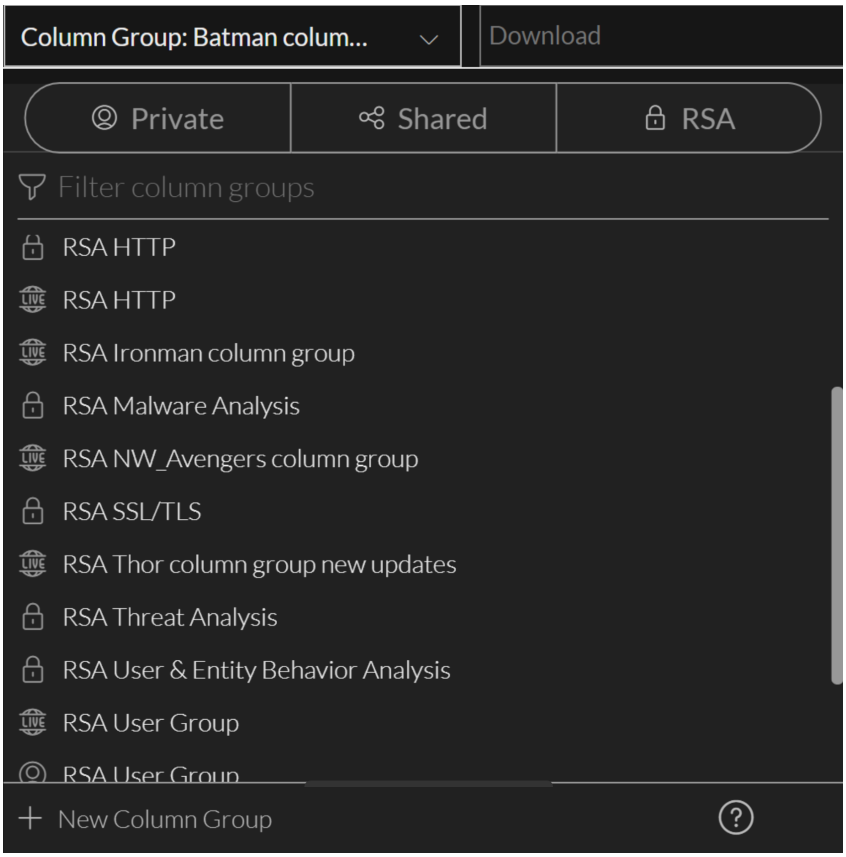


Now, if you select Decoder service and if you have not viewed the details of the Decoder recently, then the time range drop-down will not display any details as shown below:



Built-In Column Groups

NetWitness Platform has built-in column groups that include useful meta keys for specific types of investigation. The built-in groups cannot be edited or deleted, but you can create a copy of the group and edit the copy. The column groups are listed in alphabetical order in the Column Group menu in a way that makes built-in groups distinguishable from custom groups that you imported or created.

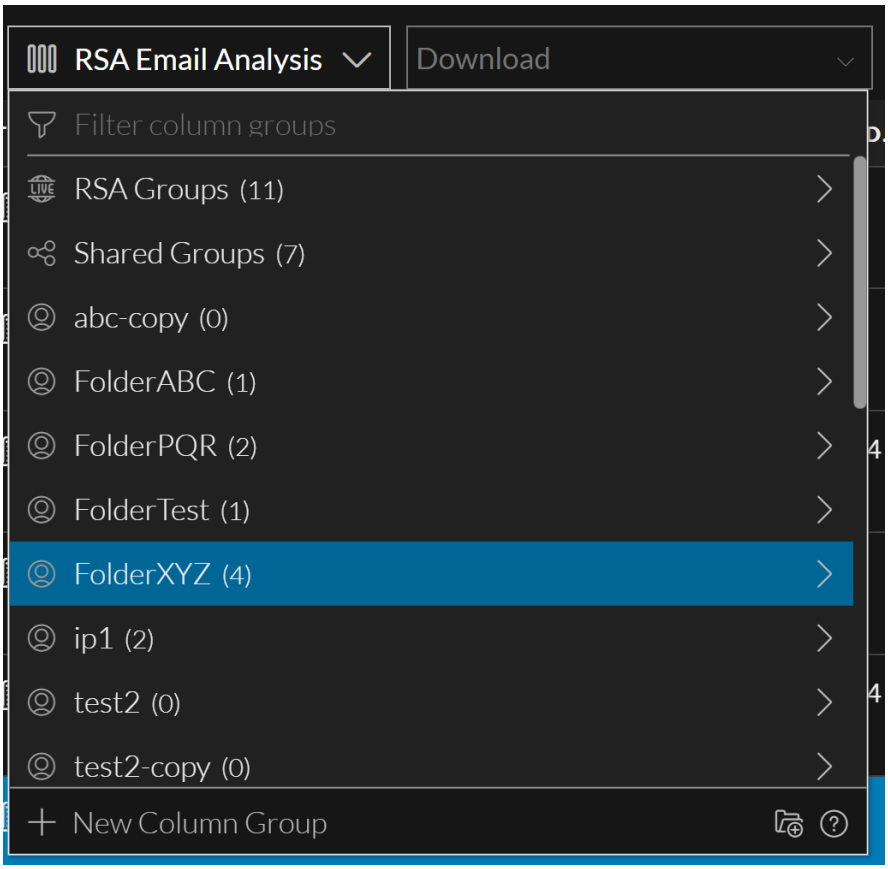


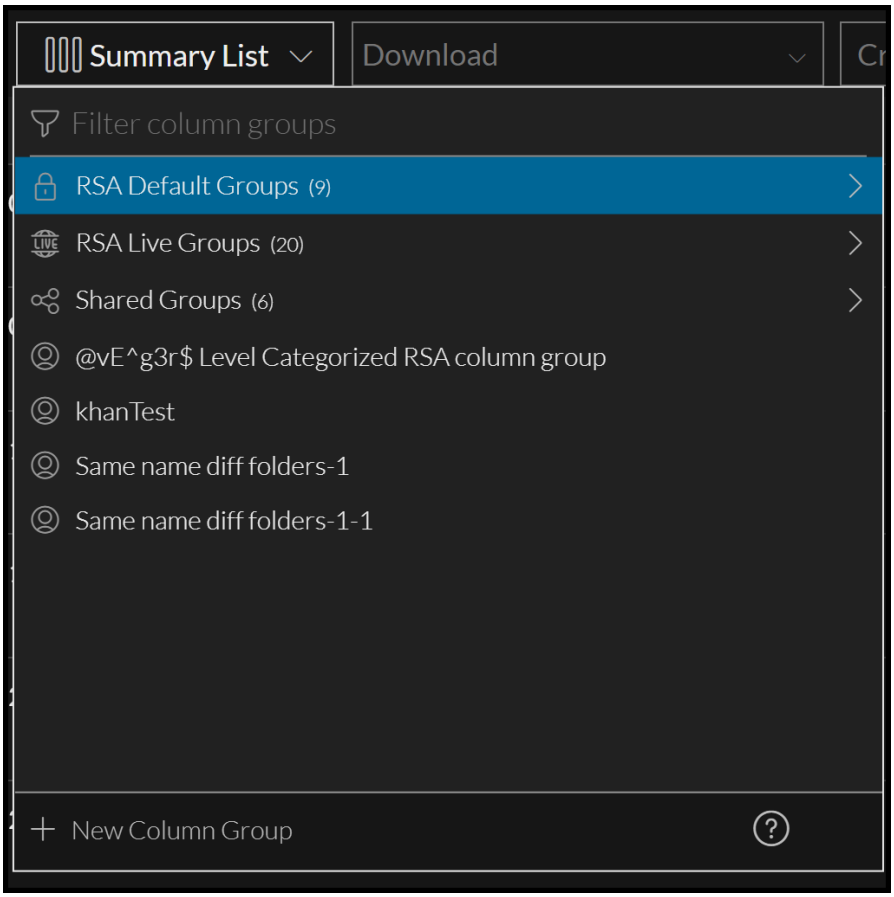
In the Legacy Events view, "RSA" precedes the name of built-in column groups. In the Events view (Version 11.4 and later), RSA precedes the name and the group is marked by the lock symbol ( ). This is an example of a selected built-in column group in the Column Groups menu. The information icon is displayed at the end of the row.

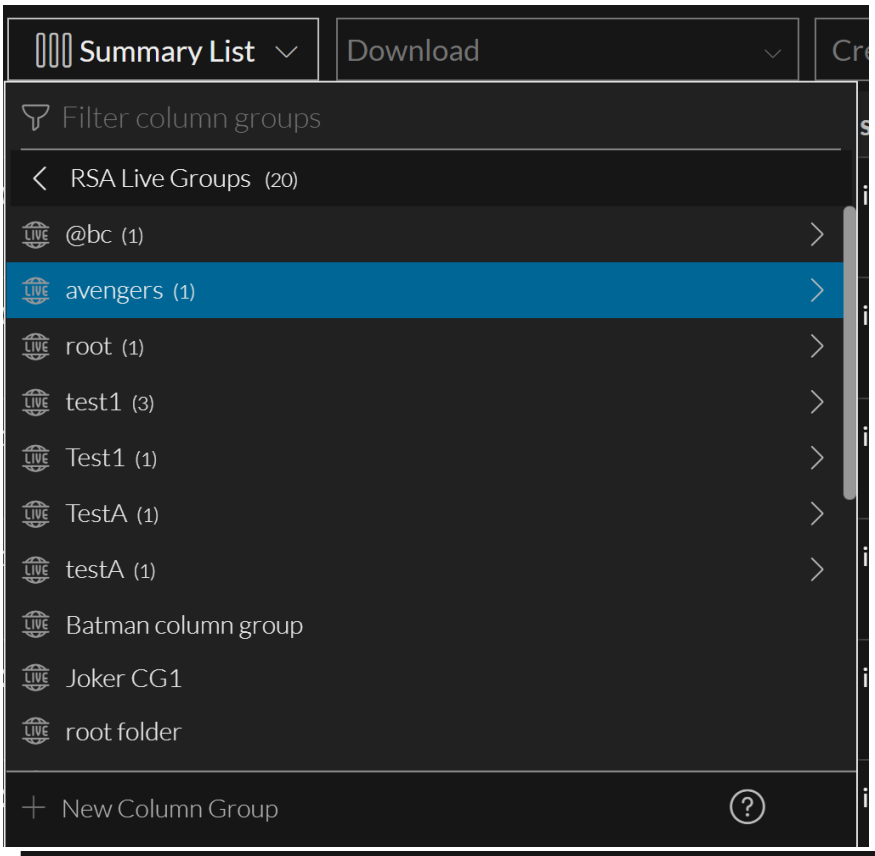


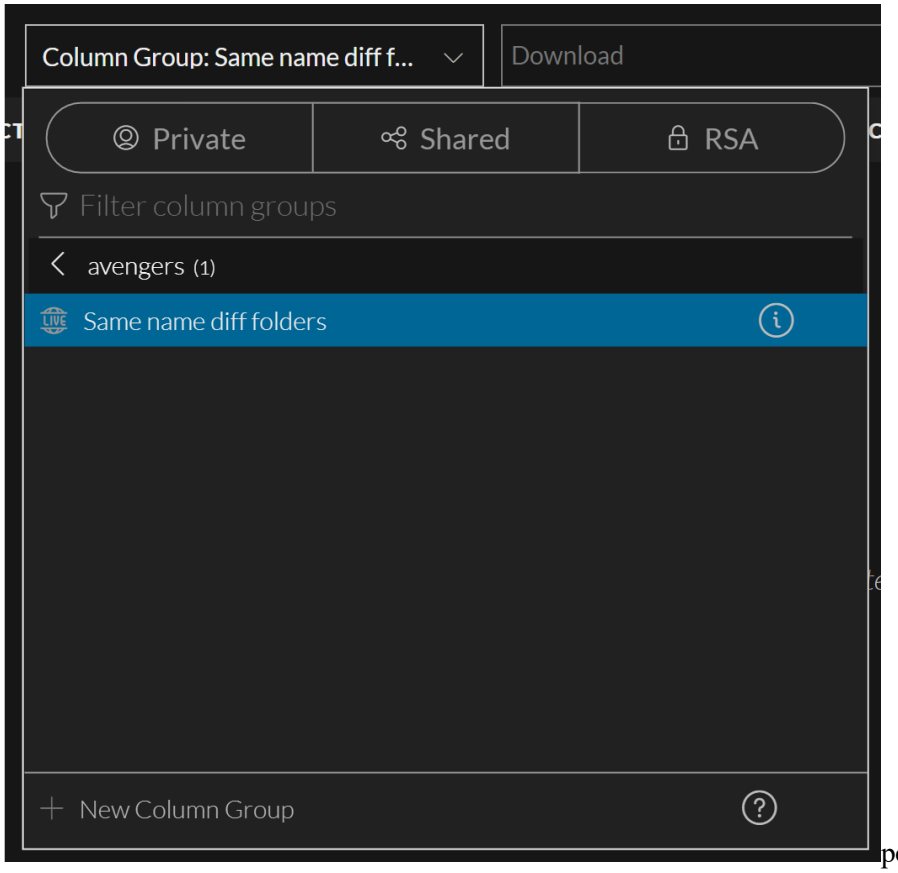
Live Column Groups

In 11.6 and later, NetWitness Platform supports deploying the investigate content from live and are marked by the live symbol () . The column groups are categorized as RSA Groups (RSA Live content and RSA OOTB Groups), and Shared Groups. The groups are displayed as non-editable folders and sub-folders except for Shared Groups that can be edited. All private content is displayed outside these groups. For example, the below image shows private content below the Shared Groups folder. The number inside () depicts the number of contents inside a folder and > symbol helps you to drill down inside the folder.


RSA Email Analysis
Download
Filter column groups
RSA Groups (11)
Shared Groups (7)
abc-copy (0)
FolderABC (1)
FolderPQR (2)
FolderTest (1)
FolderXYZ (4)
ip1 (2)
test2 (0)
test2-copy (0)
New Column Group

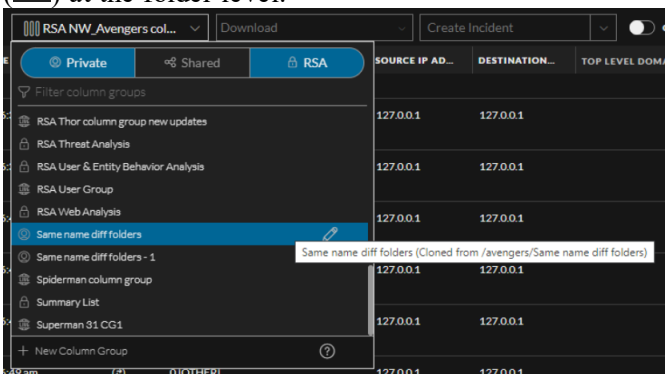

Summary List
Download
Filter column groups
RSA Default Groups (9)
RSA Live Groups (20)
Shared Groups (6)
@vE^g3r$ Level Categorized RSA column group
khanTest
Same name diff folders-1
Same name diff folders-1-1
New Column Group


Summary List
Download
Filter column groups
RSA Live Groups (20)
@bc (1)
avengers (1)
root (1)
test1 (3)
Test1 (1)
TestA (1)
testA (1)
Batman column group
Joker CG1
root folder
New Column Group


Column Group: Same name diff f...
Download
Private
Shared
RSA
Filter column groups
avengers (1)
Same name diff folders
New Column Group

pc

The column group can be copied by clicking the copy icon ( ). After copying, the copied column group is displayed under the selected location (Private folders or Shared groups). You can hover over on the cloned item to view a tooltip that displays the path from which the column group is cloned. In case you need to search for a specific column group, you can type the name of the column group in the filter field ( ) at the folder level.



These are the built-in column groups.

- **RSA Email Analysis**: Includes meta keys that are useful when investigating email-related metadata.
- **RSA Endpoint Analysis**: Includes meta keys that are useful when investigating endpoint-related metadata.
- **RSA Malware Analysis**: Includes meta keys that are useful when investigation potential malware.
- **RSA HTTP**: Includes meta keys that are useful when investigating HTTP related metadata.
- **RSA SSL/TLS**: Includes meta keys that are useful when investigating SSL/TTS analysis related metadata.
- **RSA Threat Analysis**: Includes meta keys that mark potential threats in the data set.
- **RSA User and Entity Behavior Analysis**: Includes meta keys that are useful when investigating UEBA data.
- **RSA Web Analysis**: Includes meta keys that mark anomalies in web traffic.
- **Summary List**: Includes meta keys that are useful in a general investigation. This is the default column group.

Custom Column Groups

You can create custom column groups to support scenarios that you use frequently while working in Investigate. When an administrator adds custom meta groups manually by editing the custom index file for a service, the new meta groups become available to use in column groups after the service is restarted.

Custom column groups are shared globally within your organization in Version 11.4. If you edit a shared custom column group, your changes are applied globally. If you delete a shared custom column group, the group is deleted and no longer available for all analysts. In Version 11.5 and later, you can create shared column groups as before, and can also create private column groups. When you create a group in Version 11.5, you can choose to share it or you can keep it private (default); you cannot change a shared group to private or a private group to shared.

Note: Private column groups created in the Events view are not visible or usable in the Legacy Events view.

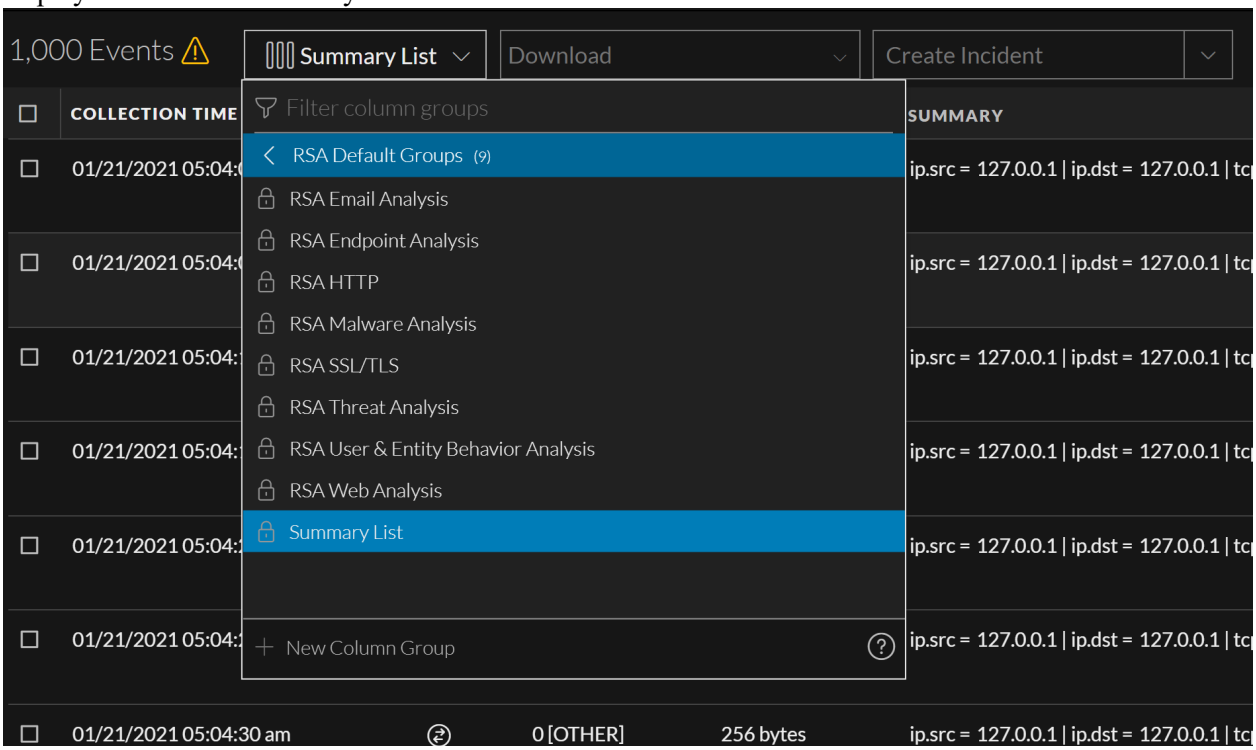
Icons identify the group type in the Column Group menu. These are examples of each type of custom column group with the edit icon displayed at the end of the row.

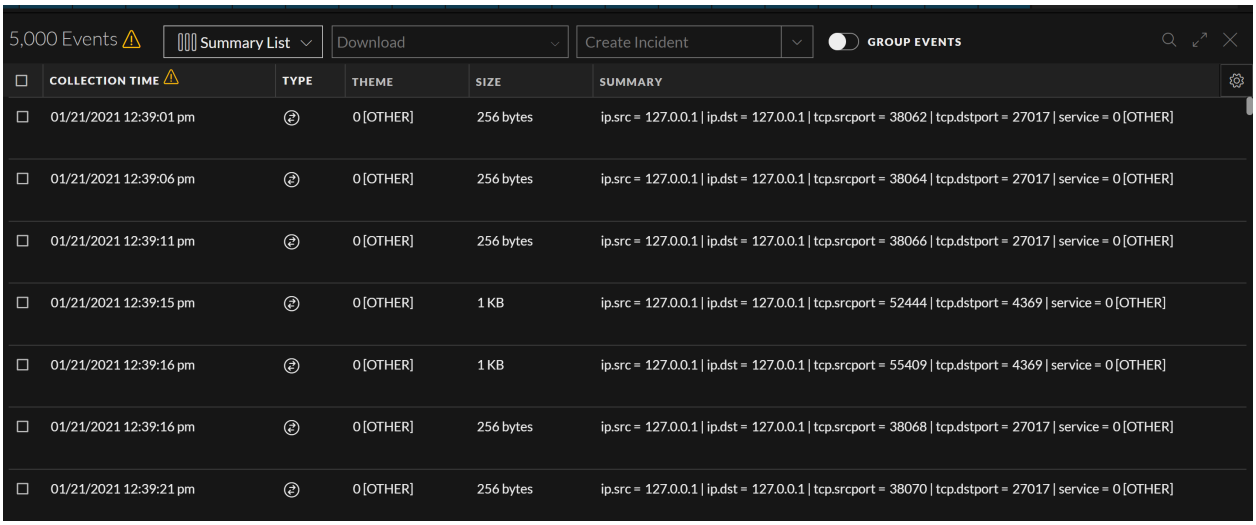


Filtering Folders

In case there are many folders, you can type the folder name and filter for a specific folder. The filtering is applicable to the current level folders and will not display folders available within a sub-folder. To search content within a sub-folder you need to navigate to the specific folder and filter.

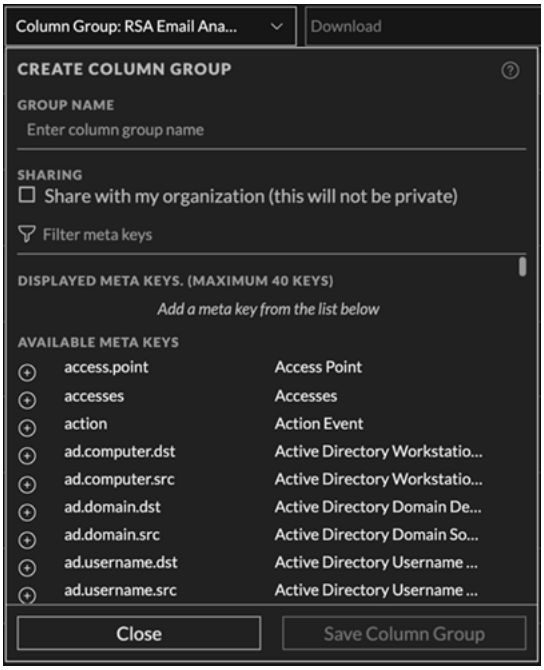
Also, when you select a specific folder, the content of the selected folder is displayed and the filter field becomes empty and when you navigate back the last selected folder is displayed. In the following example, the folder selected is RSA Groups with the its content and the column group drop-down displays the filtered Summary List folder.

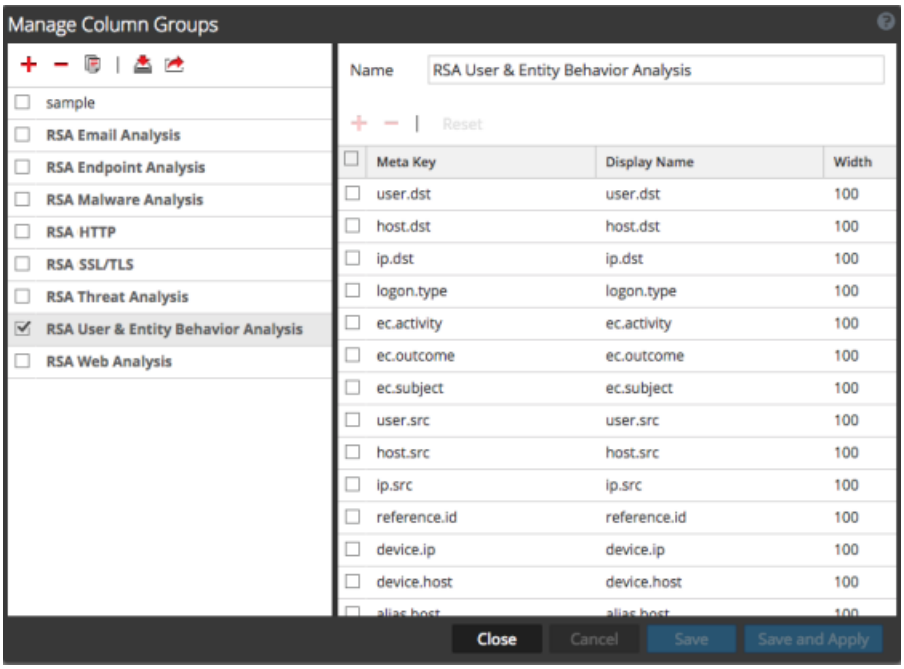




Dialogs for Managing Column Groups

While the functionality of column groups is similar in the Legacy Events view and the Events view, the user interface and some of the procedures are different. The following figures illustrate the (Events view) Create Column Group dialog and the (Legacy Events view) Manage Column Groups dialog. The Version 11.5 and later dialog includes a Sharing option.





Using options in the Create Column Group dialog and the Column Group Details dialog, you can:

- See the details of a column group.
- Create, edit, and delete custom column groups.

Using options In the Manage Column Groups dialog, you can do all of above and these additional functions:

- Clone and edit the clone of a built-in or custom column group.
- Import and export a column group.

The rest of this topic provides instructions for working with column groups in the Version 11.4 and later Events view, the 11.3 and earlier Event Analysis view, and the Legacy Events view.

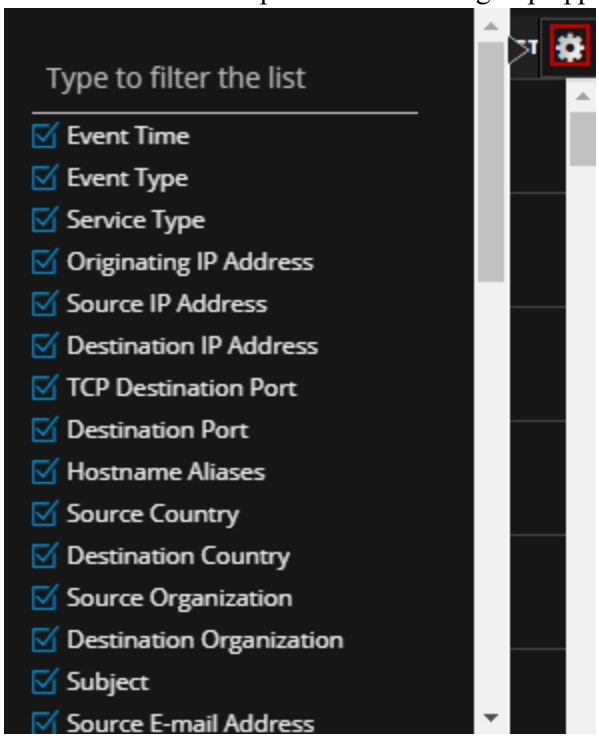
Work with Columns and Column Groups in the 11.4 and Later Events View

After the upgrade to Version 11.4, all of the existing column groups -- both built-in and custom -- are available for management in the Events view. Unless noted, the procedures in this section are for the Events view.

Manually Select Columns to Display and Adjust Column Order and Width

Note: The Column Selector was also available in the 11.3 Event Analysis view. If a column group includes a column for a meta key that your administrator has blacklisted (hidden), the data for that column cannot be displayed. The column is not available in the Column Selector and is not displayed in the Events panel.

1. With the Events list open and a column group applied, click ⚙ to display the column selector.



2. Select the meta keys or enter the name of a meta key that you want to display in additional columns.
3. Deselect the meta keys that you do not want to display in a column.
   The data is redisplayed using the selected columns.
4. To change the width of the columns in the events list, hover the cursor over the column title and drag the column divider to the right or the left.
5. To rearrange the order of the columns across the top of the events list, hover the cursor over the column title and drag the column to the right or the left.
   The changes that you make in the events list are in effect during the current session and are not retained as part of the column group. The next time the column group is applied, the original composition and order of columns is applied.

Select a Column for Sorting Events in the Events Panel (Version 11.4)

Note: You can sort events in the Events panel after results have finished loading if all connected services are updated to 11.4. or later. Sorting by column is disabled when any connected service is running an earlier version of NetWitness Platform. Version 11.4.1 has more visible sorting toggles in the column heads and the ability to view results without sorting, but otherwise it functions the same as in Version 11.4.

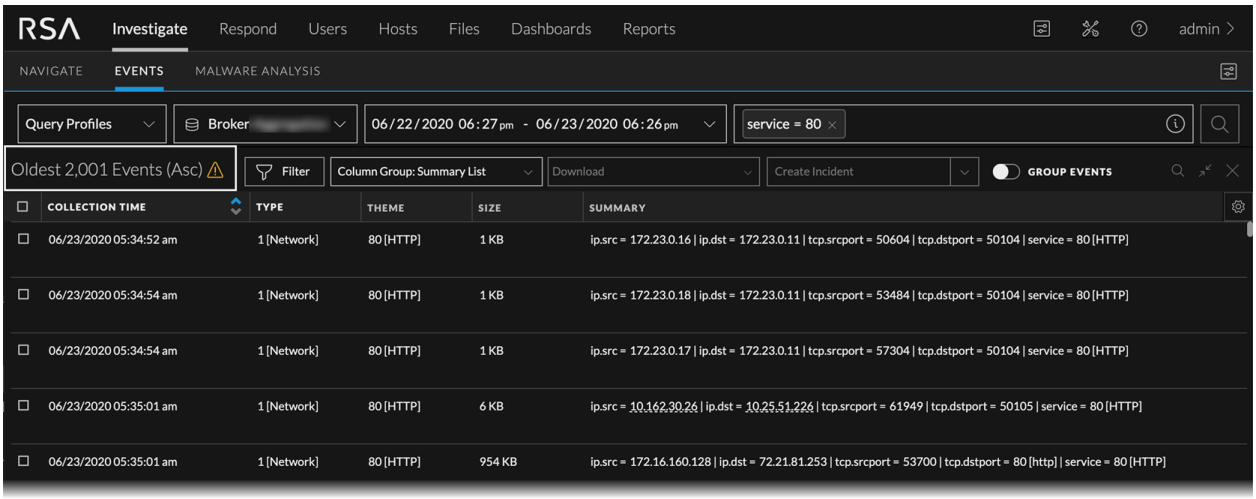
You can change the order of the events list in the Events panel based on the value for a meta key in the event. Each column title represents a meta key, and the column is populated by the values found for the meta key in the displayed events. In Version 11.4, the events in the Events panel are sorted using the method selected in the Event Preferences dialog: Ascending or Descending. If no sort method is selected, the default order is ascending (see Configure the Events View). In Version 11.4.1, the events in the Events panel are sorted only when the sort preference in the Event Preferences dialog is selected and is either Ascending or Descending. The events are not sorted if you do not have a sort preference selected under Events Preferences or if you selected Unsorted.

Sortability of a column is based on the definition of the meta key in the Broker and Concentrator index files. Columns for meta keys that are indexed by value are sortable. If the meta key is not indexed, is indexed by meta key, or has multiple values in the same event, it is not sortable.

- These are some examples of keys that are indexed by value and sortable: `time`, `eth.type`, `city.src`, `ip.src`, `ipv6.dst`, and `ipv6.src`.
- Meta entities are not sortable. For example, the meta entity `ipv6.all` is not sortable because it includes `ipv6.dst` and `ipv6.src`, and a single event has both `ipv6.dst` and `ipv6.src`.
- These are some examples of multiple value keys, which cannot be sorted: `filename`, `filetype`, and `attachment`. A single event can have more than one file and therefore more than one value for `filename`, `filetype`, and `attachment`.
- These are some examples of meta keys that cannot be sorted because they are not indexed or not indexed at the values level: `password`, `query`, and `size`.

Sorting by Column (Version 11.4.1 and Later)

The initial view of the Events list with the sorting preference set to Unsorted and no column sorting has an event count in the title, with no indication of a sorting method applied to a column. If the event sorting preference is set to Ascending, the count label is "Oldest 1,000 Events." If the event sorting preference is set to Descending, the count label is "Newest 1,000 Events." In the figure below, the Ascending method is in effect, more than 2001 events matched the query, and only the oldest 2001 are displayed. Clicking the amber warning triangle displays an explanation. Refer to Configure the Events View for more information about the sorting preference.



When you move the mouse over a column title, sortable columns have a pair of arrows after the column title, one pointing up for ascending and one pointing down for descending ( ). You can choose one sort column and the direction of the sort. A blue up arrow ( ) indicates that ascending sort order is in effect; which means the earliest events or the lowest numbers, or the text strings beginning with an 'A' appear first. A blue down arrow ( ) indicates that descending sort order is in effect; which means the latest events or the highest numbers, or the text strings beginning with a 'Z' appear first.

- When a column has a blue arrow, you can click the white arrow to change the sort order. When you change the sort order, a blue progress bar is displayed in the Events list title bar to show progress. As sorting begins, there is a short segment on the left side of the window; as sorting progresses the blue color extends to the right across the entire title bar. The directional arrow does not change until the events are re-sorted in the chosen sort order.
- To change the column to unsorted, you can click the blue arrow. Both arrows are white now to show that the column is unsorted. This figure shows the Type column sorted in ascending order.

- If a column is not sortable, no arrow is displayed when you hover the mouse over the column title. Instead a tooltip explains why it is not sortable.

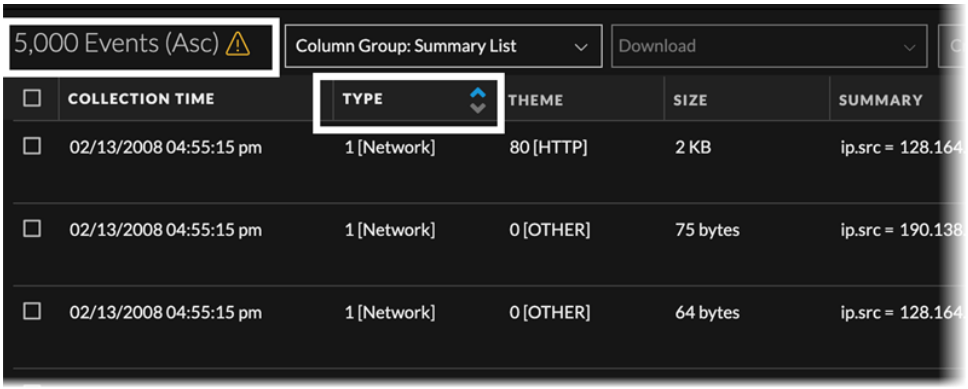
Sorting on a column is done on the client side without re-executing the query if the number of displayed results is less than the events limit set by the administrator. If there are more results that are not displayed because the number of results exceeded the events limit, a new query is submitted with the new sort order, and the same service, time range, and filters. The current results are removed, a spinner indicates progress, the Cancel button becomes available, the reconstruction closes, and progress is visible in the Query console.

Note: The re-sorting of events takes place in the browser when the number of results of the original query is less than the event display threshold.

To change the sort order or the sort column:

1. Move the mouse over the column titles to find a sortable column.
   If a column is not sortable, a tooltip that explains the reason is displayed.

2. To sort the list based on a column, move the mouse over a sortable column and click one of the arrows ( ).
   The arrow turns blue and the events are reloaded in the selected order. If both arrows are white, the column is not being used to sort the events list. If one arrow is blue, the column is being used to sort the events list, and the sort order (Asc or Desc) is appended to the events count in the title bar. This figure shows a column sorted in ascending order. When a column is descending order, (Desc) is appended to the event count.



   a. Click a white arrow to sort the events list in that order.

   b. Click a blue arrow to return to unsorted order.

Sorting by Column (Version 11.4)

When you move the mouse over a column title, sortable columns have an up or down arrow (or) after the column title. You can choose one sort column and the direction of the sort. An up arrow indicates that Ascending sort order is in effect; which means the earliest events or the lowest numbers, or the text strings beginning with an 'A' appear first. A down arrow indicates that Descending sort order is in effect; which means the latest events or the highest numbers, or the text strings beginning with a 'Z' appear first. When you select a sort column, it is sorted in descending order by default, with events having a null value for the meta key first.

- A column that is being used to sort the events list has a bright white arrow indicating the direction that you can choose for sorting: click to change to Ascending or to change to Descending order. When you click to change to Ascending sort order, the directional arrow does not change until the events are re-sorted in ascending order. The same behavior applies when you click the to change to Descending order.
- If a sortable column is not being used to sort the events list, the arrow is dimmed. If a column is not sortable, no arrow is displayed when you hover the mouse over the column title. Instead a tooltip explains why it is not sortable.
- If you click the arrow on a different column, the column is sorted in the same order as the previously active sort column. You can select a different sort order if desired.

Sorting on a column is done on the client side without re-executing the query if the number of displayed results is less than the events limit set by the administrator. If there are more results that are not displayed because the number of results exceeded the events limit, a new query is submitted with the new sort order, and the same service, time range, and filters. The current results are removed, a spinner indicates progress, the Cancel button becomes available, the reconstruction closes, and progress is visible in the Query console.

Note: The re-sorting of events takes place in the browser when the number of results of the original query is less than the event display threshold. If some of those events have the exact same time, they will not change order as you might expect when you reverse the sort order.

To change the sort order or the sort column:

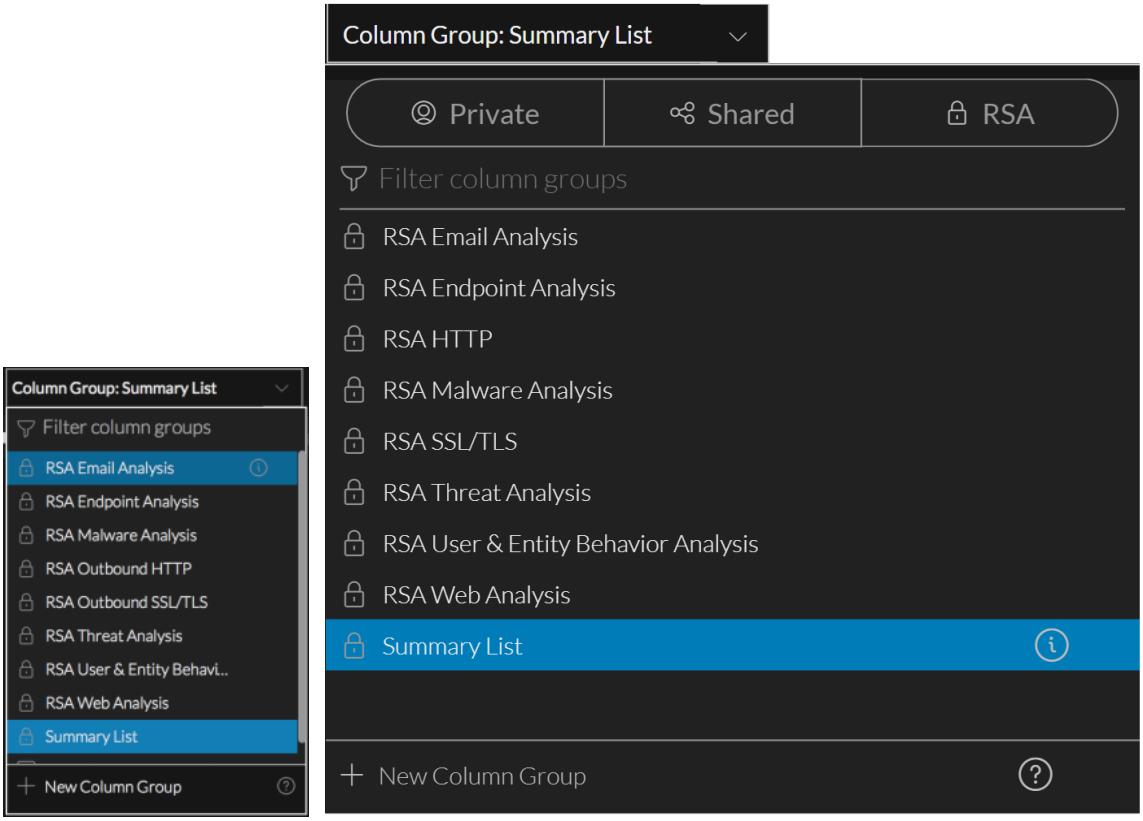
1. Move the mouse over the column titles to find a sortable column.
   If a column is not sortable, a tooltip that explains the reason is displayed.
2. To sort the list based on a column:
   a. Move the mouse over a sortable column and click the arrow ( or ).
   The events are sorted in the correct sort order. If you hover over the column title, you can see that the arrow is no longer dimmed. A column that is being used to sort the events list has a bright white arrow that you can click to change the sorting direction.
   b. To change the sort order, click  to change to Ascending or  to change to Descending order.
   The direction of the arrow changes and the events are reloaded in the selected order.

View the Meta Keys Included in a Column Group

To view details of a column group:

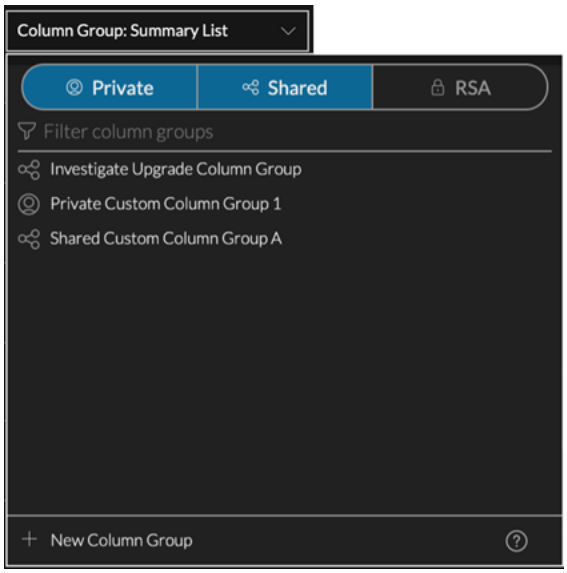
1. Go to **Investigate > Events** and click  to load events.
   The events for the default service and the default time range are loaded in the Events panel. The Summary List column group or the column group from your last session is applied to the list.
2. To display the Column Groups menu, click the Column Groups menu title. The Column Group menu title includes the title of the currently selected column group. If this is your first visit after logging in, the Summary List group is selected; any subsequent visits use the column group selected in the previous session. When opened, the menu displays a list a list of built-in column groups (RSA), shared custom column groups, and your private custom column groups. Visibility options (Version 11.5) and a filter field make it easier to find a particular column group. The figure on the left shows the Version 11.4 menu initially when Summary List is selected by default and the first column group in the list is highlighted. The figure on the right shows the Version 11.6 menu initially when Summary

List is selected by default and all types of column groups are visible: Private, Shared, and RSA.

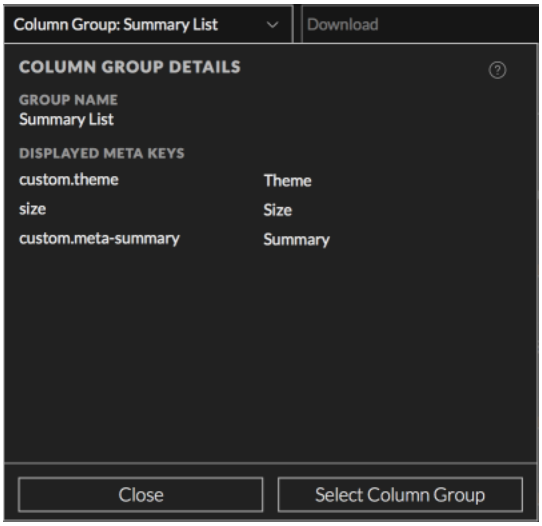


3. (Optional) To control the types of column groups that are visible in the list, use any combination of the visibility options (blue = selected, black = not selected):
   Private = display private groups that only you can manage
   Shared = display shared groups that anyone in your organization can manage
   RSA = display built-in groups that only RSA can manage
   The visibility options work together with the Filter Column Groups field. If the visibility option is hiding built-in groups (which include "RSA" in the group name) and you search for a name that contains "RSA," the list is empty. The figure below shows private and shared visibility options

selected.



4. Hover over the **Summary List** group and click the information icon ( ) to see which columns are included in the group.
   This figure shows the columns for the Summary List. The Collection Time and Type column are always the first two columns in the Events list, but are not listed in the Column Group Details dialog.

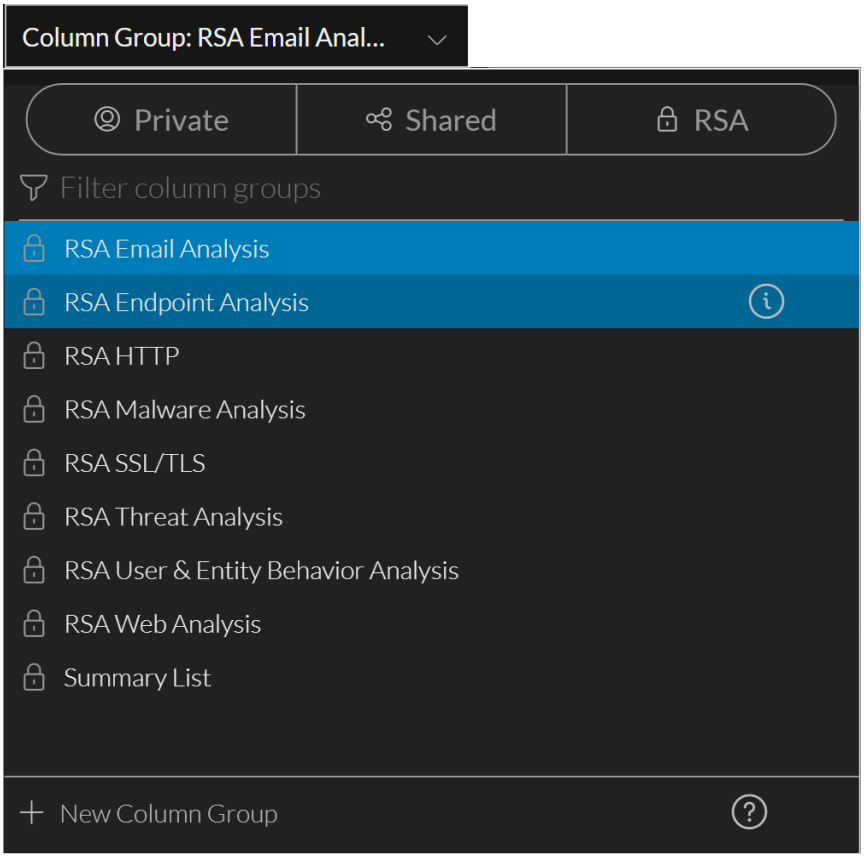


5. Do one of the following.
   a. To close the dialog, click **Close**.
   b. If you want to apply the column group, click **Select Column Group**.
      The dialog closes and the Events list is updated to reflect the selected column group.

Select a Column Group

1. With the Events panel open in the 11.4 or later Events view, click the **Column Group** menu title.
   The menu drops down to display a list of column groups with a filtering option and a New Column

Group option. The list is sorted alphabetically and the selected column group is displayed in the menu label. The first option in the list is highlighted. The selected column group has a slightly different background color than the highlighted column group.
The following figure shows the menu after RSA Endpoint Analysis was highlighted, but RSA Email Analysis is still selected.

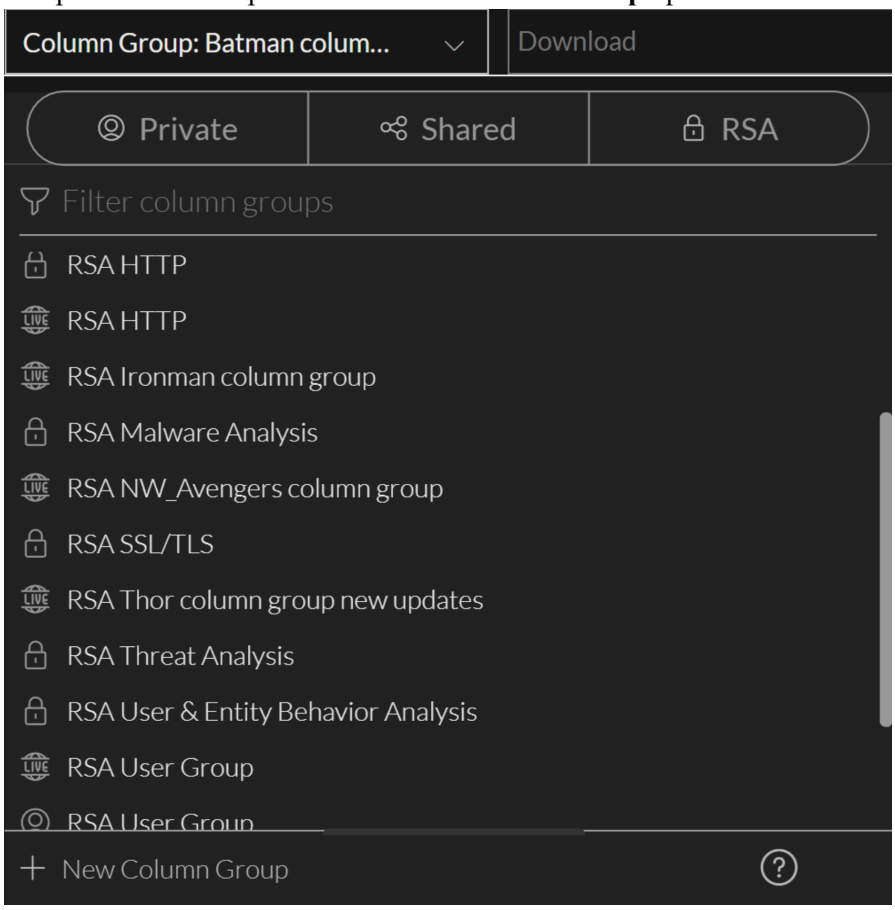


2. Do one of the following:
   a. If the highlighted group is the one you want to apply, press **ENTER**.
   b. (Version 11.5 and later) If you want to see only certain types of groups, use the visibility options (**Private**, **Shared**, and **RSA**) to hide one or two group types.
   c. Begin typing text in the **Filter column groups** field to search for a column group name. As you type, the list is filtered to show only the column group names that contain that string.
   When you see the group that you want to apply, click it or use the down or up arrow to highlight it, then press **ENTER**.
   The Events list is refreshed to include only columns in the selected column group, and the menu title includes the selected group name. Your selection persists when you navigate away from the Events view. The order of the columns in the Events list reflects the order of the meta keys in the column group. A column group may contain more columns that are only visible when you scroll to the right. For optimal viewing, the first 15 columns are displayed by default when you select a column group.

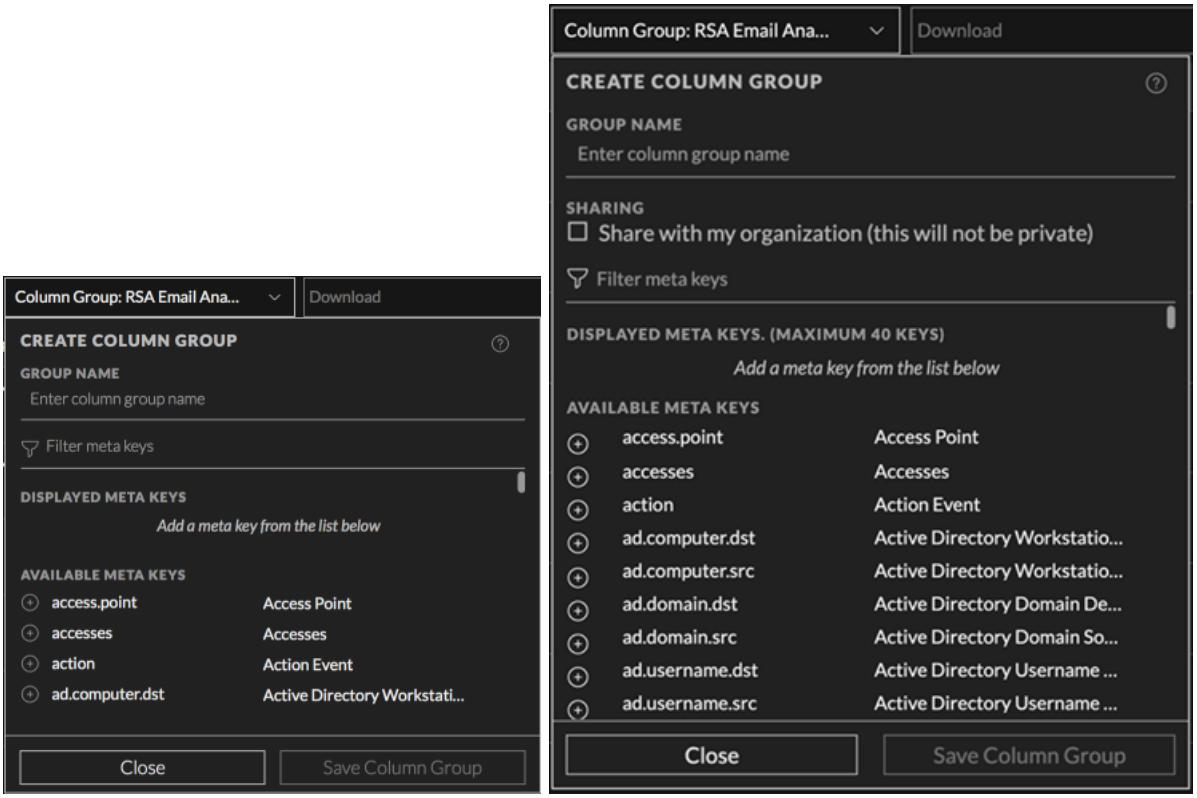
Note: If a meta key in a column group is not part of the selected service, it does not appear in the Filter Events panel or in the Events panel.

Create a Custom Column Group

1. Go to **Investigate > Events** and submit a query to load data in the Events panel.
2. In the Events panel toolbar, click the **Column Group** menu title.
   The menu drops down to display a list of column groups with the Visibility Options and Filter Column Groups field at the top and the + **New Column Group** option at the bottom.

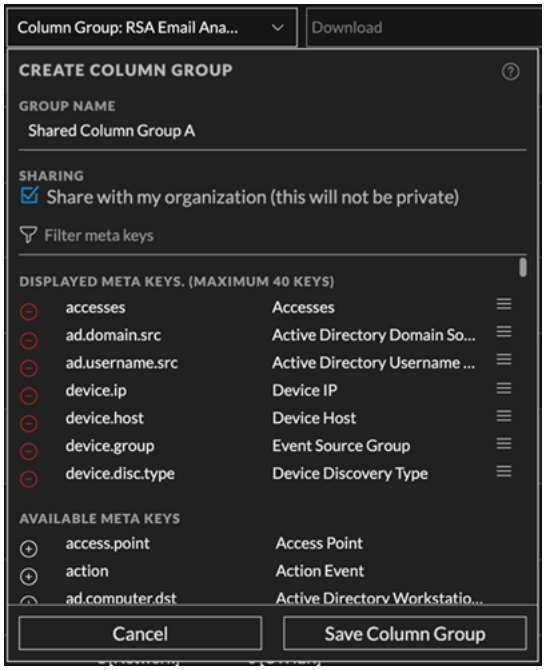


3. Select + **New Column Group**.
   The Create Column Group dialog is displayed. Version 11.5 includes the Sharing option.



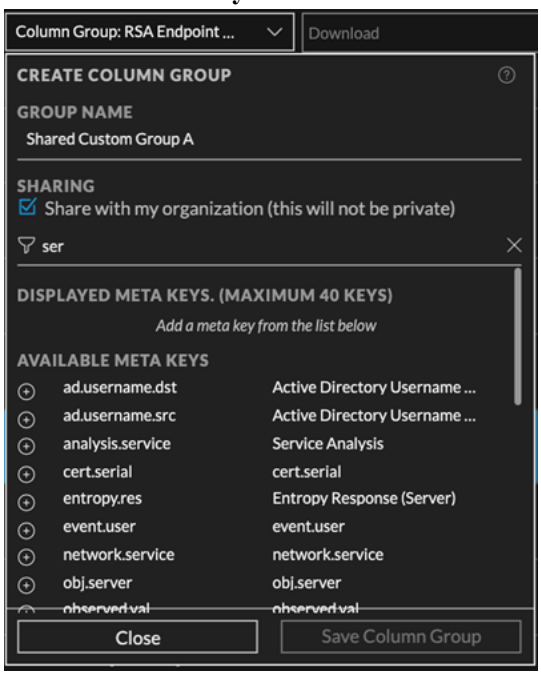
4. In the **Group Name** field, type a unique name (maximum length of 256 characters) for the new column group, for example, **Custom Column Group A**.

5. (Version 11.5 and later) If you want to share the new column group with your organization, set the **Share with my organization** option.



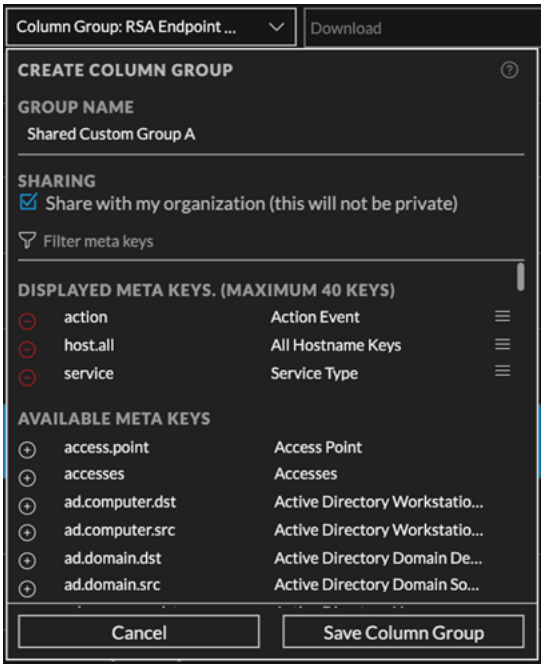
6. To add a meta key to the column group, select and add each meta key as follows:

   a. Type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Available Meta Keys** list.



   b. When you see the meta key that you want to add, click the add icon  that precedes the meta key name.
   The meta key is added to the end of the Displayed Meta Keys list. (This list is also filtered using

the text you typed.) The maximum number of meta keys in a column group is 40. If you attempt to add another meta key when 40 are already included in the Displayed Meta Keys list, a message advises you that the group has the maximum number of meta keys.



7. (Optional) To find and remove a meta key from the column group, type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Displayed Meta Keys** list. When you see the column that you want to remove, click the remove icon ( ) that precedes the meta key name in the **Displayed Meta Keys** list.
   The meta key is moved back to the Available Meta Keys list.
8. (Optional) To change the order of the displayed meta keys in the Displayed Meta Keys list, place the cursor over the list order icon ( ). When the cursor changes to the drag and drop icon ( ), drag the meta key up or down in the list.
9. Do one of the following:
   a. To close the dialog without creating the custom column group, click **Cancel**.
   b. To create the group, click **Save Column Group**.
      The new column group is saved and becomes available for all analysts. The buttons change to Done and Select Column Group.
10. Do one of the following:
    a. To close the dialog, click **Done**.
    b. To close the dialog and select the new column group, click **Select Column Group**.
       The new group is added to the Column Groups menu (in alphabetical order), and if you clicked Select Column Group, the Events list is updated to show the columns in the new column group.

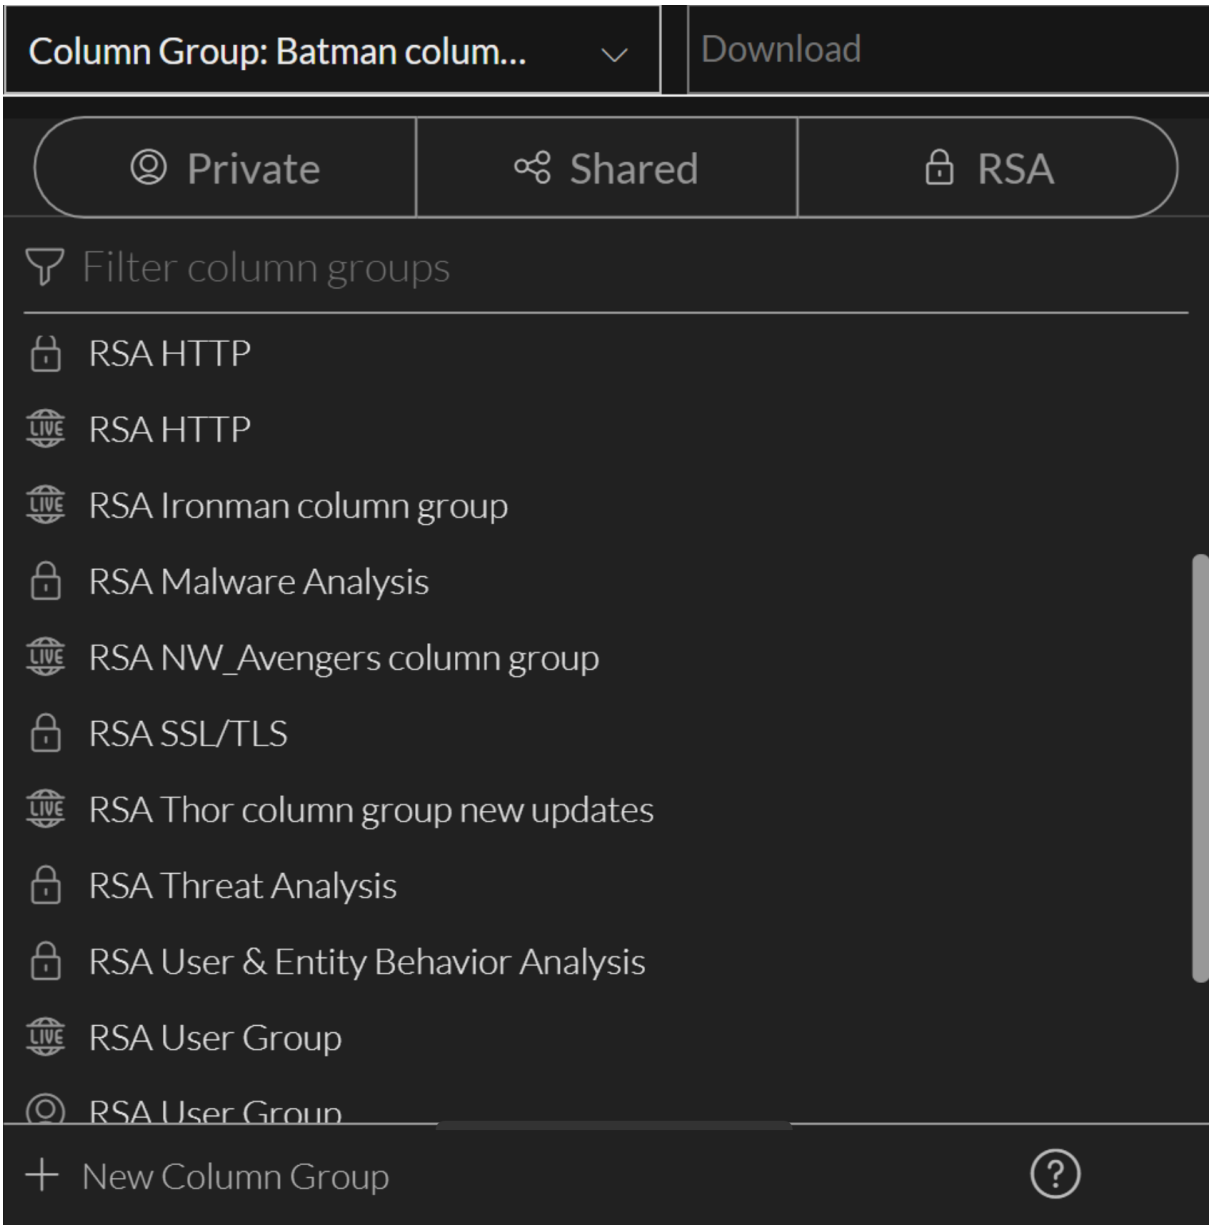
Delete a Custom Column Group

You can delete any custom column group that is not currently applied in the Events list and not part of a query profile. The built-in column groups are read only, and cannot be deleted. In Version 11.5 and later, a confirmation message allows you to confirm or cancel the deletion. When you delete a custom column group, it is removed from the Column Group menu.

Caution: When you delete a custom column group (Version 11.4) or a shared column group (Version 11.5), the effect is global and the group is no longer available to any analyst.

To delete a custom column group:

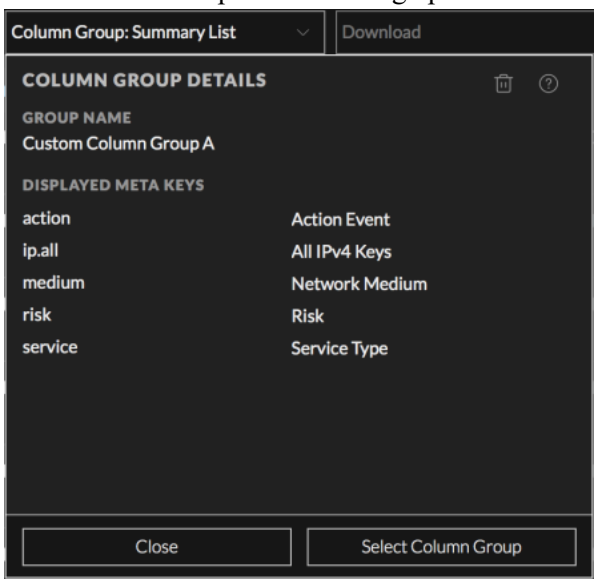
1. Go to **Investigate > Events** and click  to load events.
   The events for the default service and the default time range are loaded in the Events panel. The Summary List column group or the column group from your last session is applied to the list. This figure shows the initial view with the Summary List column group selected. The label on the Column Group menu includes the name of the selected column group.



2. To delete a column group, highlight a custom column group as shown in the following figure and click the edit icon ( ) to the right of the name.



3. The Column Group Details dialog opens with the details for the selected group displayed.



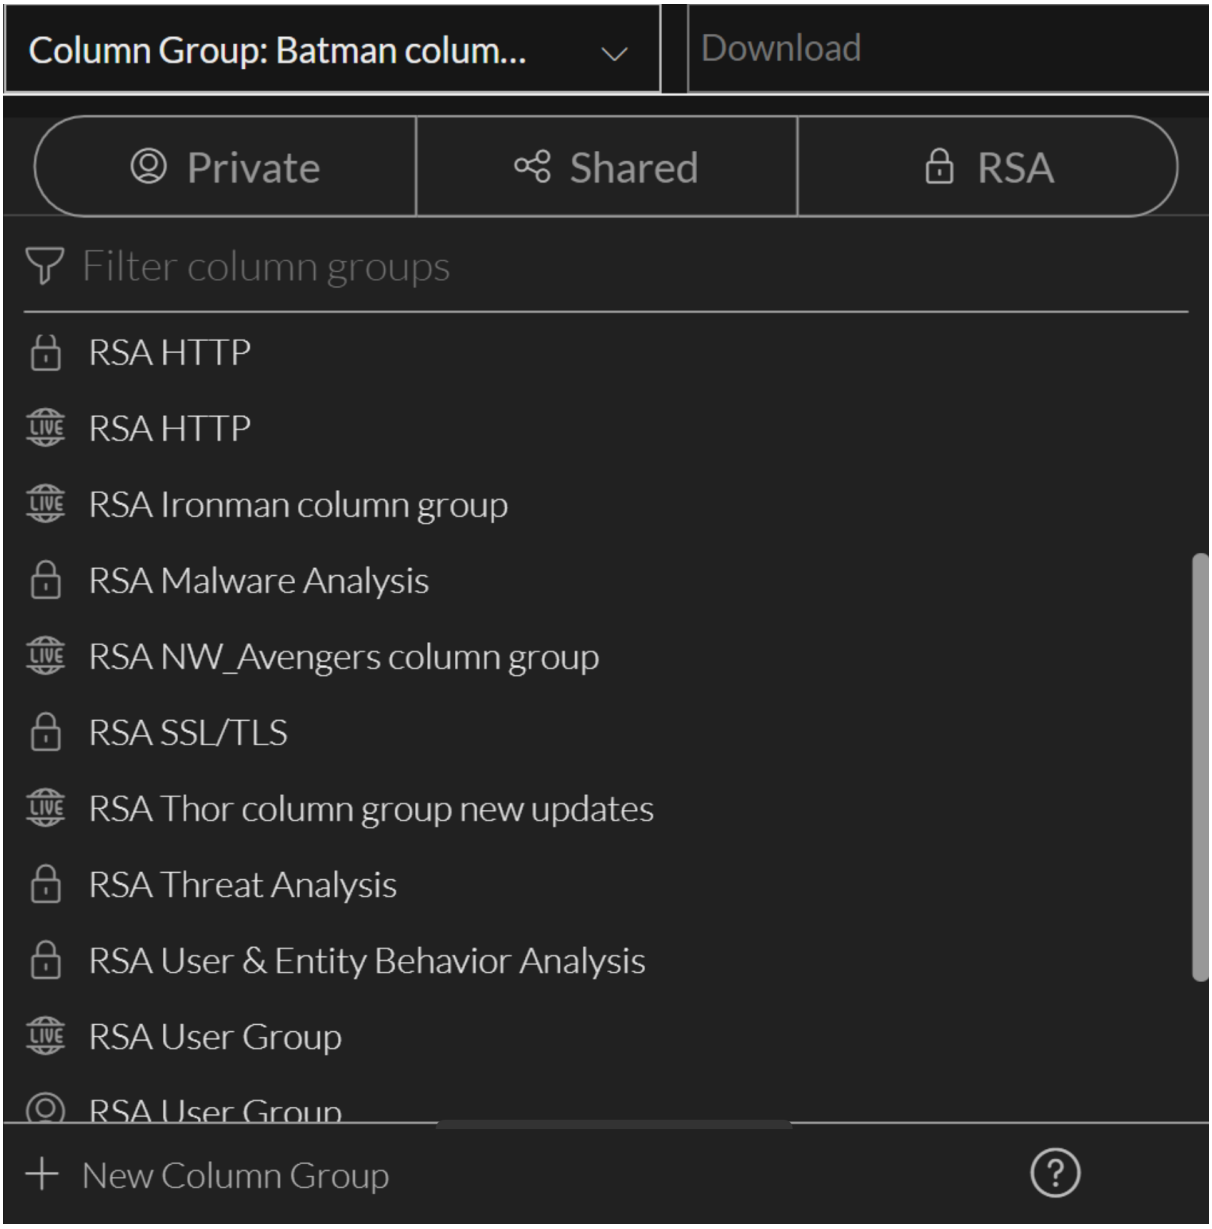
4. Click the delete group icon ( ).
   If the column group is currently in effect, the following message is displayed: `This column group cannot be deleted because it is currently active.`
   In Version 11.5, a confirmation message gives you the opportunity to confirm or cancel the deletion. Click **Cancel** or **Delete Column Group**.
   In Version 11.4, if the column group is not in effect and is not a built-in column group, there is no request for confirmation before the column is deleted.
   The group is deleted and removed from the Column Groups menu. The column group no longer appears anywhere for any analyst working in Investigate.

Edit a Custom Column Group

You can create a shared or private copy of any column group that is not open for editing. After you create the copy, you can edit the new group in the usual way.

1. Go to **Investigate> Events** and submit a query to load data in the Events panel.

2. In the Events panel toolbar, click the **Column Group** menu title.
   The menu drops down to display a list of column groups.

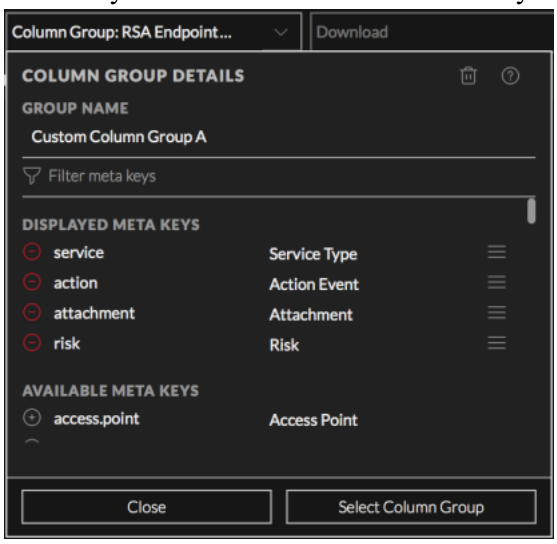


3. Highlight the column group that you want edit. This figure shows a custom column group highlighted, with the edit icon displayed to the right.

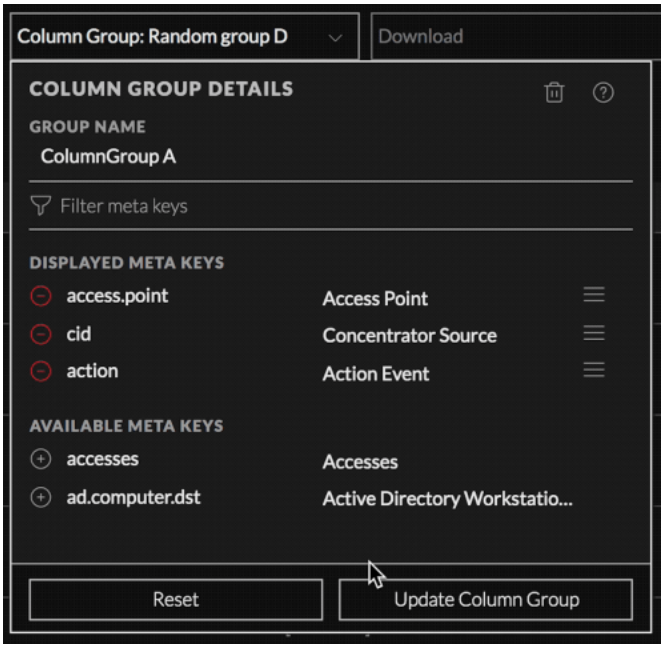


4. Click the edit icon ( ).
   The Column Group Details dialog is displayed so that you can edit the Group Name and Displayed

Meta Keys. You can add or delete meta keys and rearrange the order of the meta keys in the list.



5. (Optional) In the **Group Name** field, edit the name of the column group.
6. (Optional) To add a meta key to the column group, select and add each meta key as follows:
   a. Type a text string in the **Filter meta keys** field and look for meta keys that contain that text in the **Available Meta Keys** list. Or just scroll through the list to find the meta key.
   b. When you see the meta key that you want to add, click the add icon  that precedes the meta key name.
   The meta key is added to the end of the Displayed Meta Keys list. (This list is also filtered using the text you typed.) This figure shows the group name changed to Column Group C and `access.point` added to the Displayed Meta Keys list.



7. (Optional) To find and remove a meta key from the column group, type a text string in the **Filter meta keys** field to look for meta keys that contain that text in the **Displayed Meta Keys** list, or simply scroll through the list. When you see the column that you want to remove, click the remove icon ( ) that precedes the meta key name in the **Displayed Meta Keys** list.
   The meta key is moved back to the Available Meta Keys list.
8. (Optional) To change the order of the displayed meta keys in the Displayed Meta Keys list, place the cursor over the list order icon ( ). When the cursor changes to the drag and drop icon ( ), drag the meta key up or down in the list.
9. Do one of the following:
   a. To close the dialog without saving the changes to the custom column group, click **Reset**.
   b. To save the edits to the column group, click **Update Column Group**.
      The updated column group is saved globally for all analysts, and the buttons change to Done and Select Column Group.
10. Do one of the following:
    a. To close the dialog, click **Close**.
    b. To close the dialog and select the updated column group, click **Select Column Group**.
       The column group is updated, and if you clicked Select Column Group, the Events list is updated to show the columns in the new column group.

Create a Copy of a Column Group (Version 11.5 and Later)

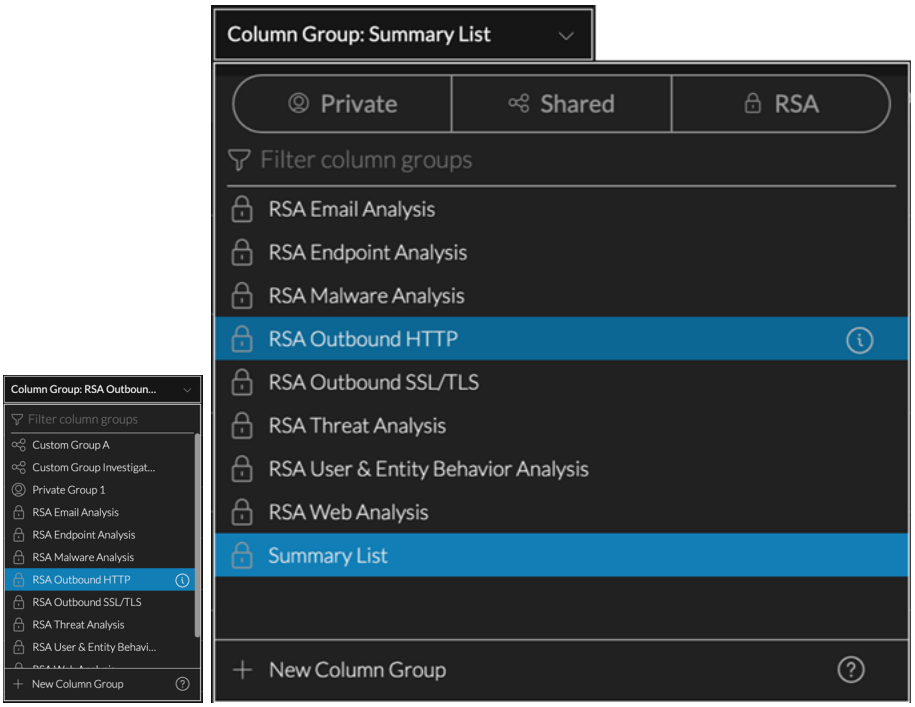
You can copy any column group, built-in or custom, shared or private, as long as it does not have unsaved edits in progress. This is useful when you want a customized version of a built-in group. Also since you cannot change a custom group from private to shared or from shared to private, creating a copy allows you to select a different Sharing setting. When you create a copy of a column group, the same name is used with a number appended. For example, if you copy RSA HTTP, the first copy is named RSA HTTP-1, and a second copy of the same group is named RSA HTTP-2. After you create the copy, you can edit the new group to give it a new name and manage meta keys in the group.

Note: Some column groups created in the Legacy Events view may have more the 40 columns, which is above the limit for column groups in the Events view. If you copy a group with more than 40 columns, you must remove the excess columns when you edit the column group.

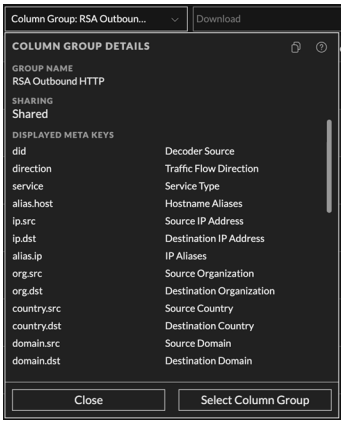
To copy a column group:

1. Go to **Investigate> Events** and submit a query to load data in the Events panel.
2. In the Events panel toolbar, click the **Column Group** menu title.
   The menu drops down to display a list of column groups with the Filter Column Groups field at the top and the + New Column Group option at the bottom. The first group on the list is highlighted, and the selected group has a light blue background.
3. Highlight the column group that you want copy. This figure shows RSA HTTP highlighted. The information icon ( ) is displayed to the right. On the left is the Version 11.4 menu and on the right is

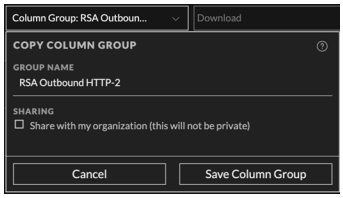
the Version 11.6 and later menu.



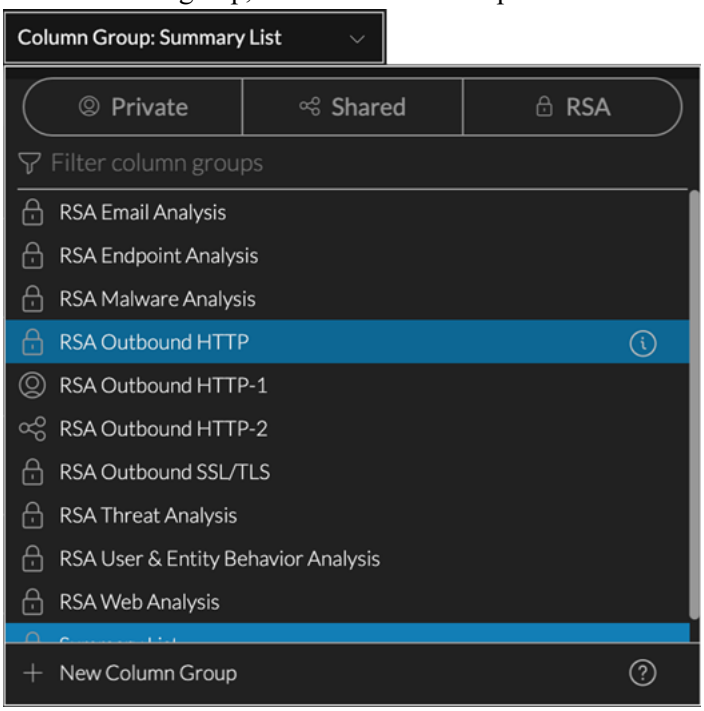
4. Do one of the following:
   a. Click the information icon ( ).
   b. Click the edit icon ( ).
   The Column Group Details dialog is displayed. This figure shows the dialog for a built-in group.



5. Click the Copy icon ( ).
   The Copy Column Group dialog is displayed with a `-n` appended to the column group name. The following figure has -2 because it is the second copy of this column group.

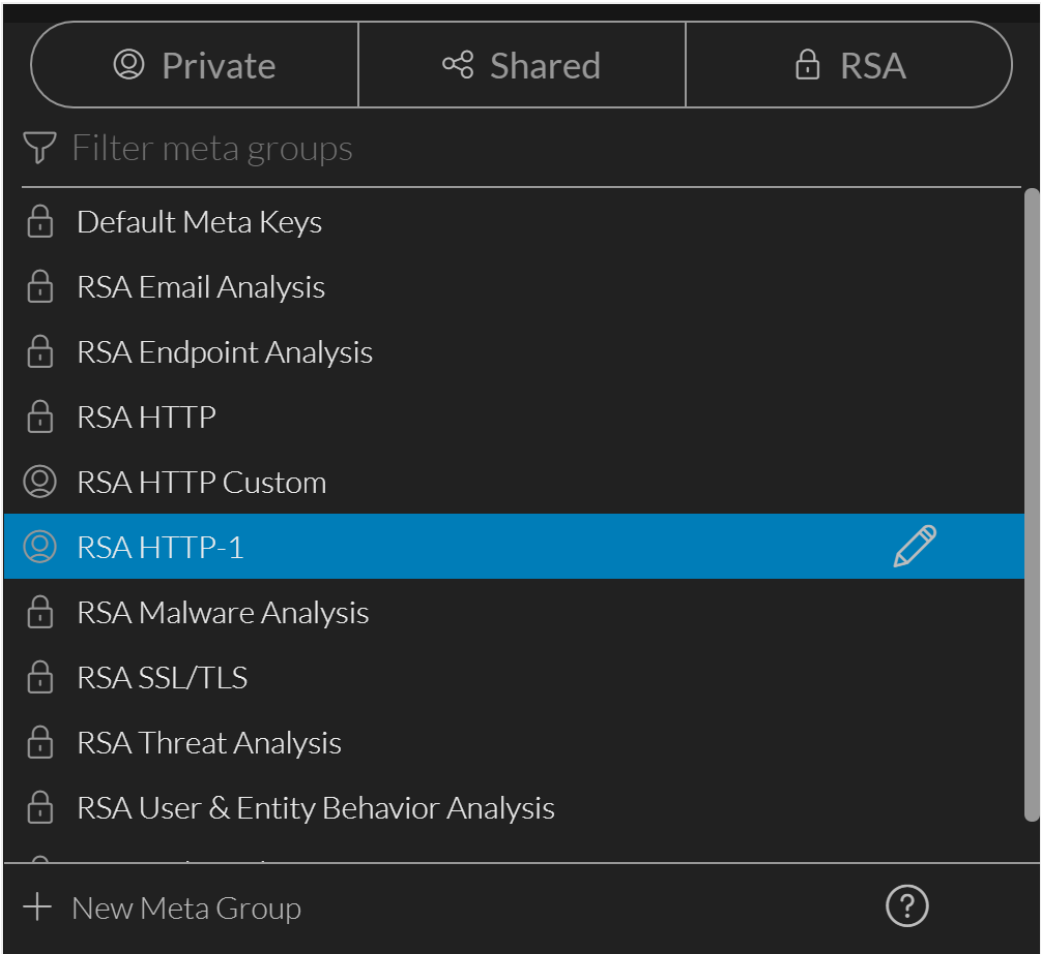


6. (Optional) In the **Group Name** field, edit the name of the column group.
7. If you want to share the new column group with your organization, set the **Share with my organization** option. By default the new group is private.
8. Do one of the following:
   a. To close the dialog without copying the group, click **Cancel**.
   b. To save the copy of the column group, click **Save Column Group**.
   The copy of the column group is saved, and the buttons change to Done and Select Column Group.
9. Do one of the following:
   a. To close the dialog, click **Close**.
   b. To close the dialog and select the copy of the column group, click **Select Column Group**.
   The column group is copied, and if you clicked Select Column Group, the Events list is updated to show the columns in the copy of the column group. The figure below has two copies of the RSA HTTP column group, one shared and one private.



1. Do one of the following:
   a. To close the dialog without editing, click **Close**.
   b. To close the dialog and select the copy of the meta group, click **Select Meta Group**.
   The group is added to the Meta Group menu. The figure below has a private copy of the RSA

HTTP meta group.



Create a Column Group Folder

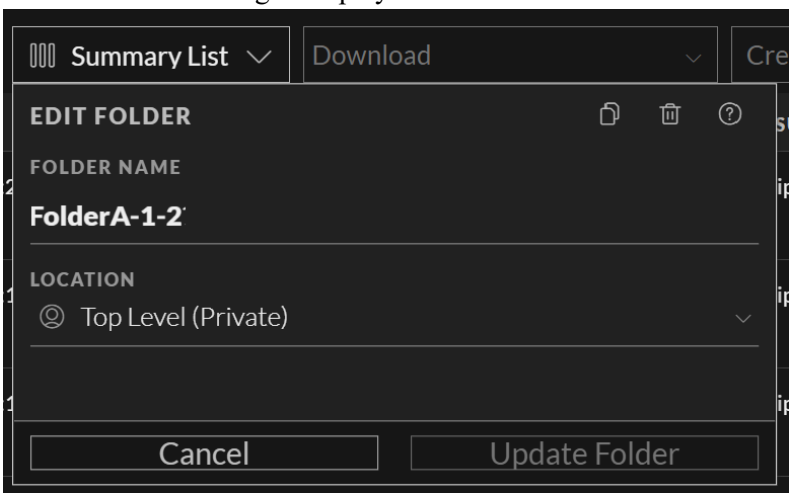
You can create custom column group folders which reside at the current level and are be added as a private or shared folder. And, if the folder name already exists then you are prompted to provide a unique name.

1. With the Filter Events panel open in the Events view, click the Column Groups menu title. The menu drops down to display a list of column groups and folders.
2. Click  .
   The Create Folder dialog is displayed.
3. In the **Folder Name** field, type a unique name (maximum length of 255 characters) for the new meta group folder.
4. Click **Create Folder**.

Edit and Move Column Group Folder

After you create a column group folder you can edit or move it, however the folders inside RSA Groups (RSA Live content and RSA OOTB Groups) cannot be edited and moved. The folders inside private and shared folders can be edited and moved only within their respective groups. For example, you cannot move a shared folder into a private folder and vice-versa.

1. With the Filter Events panel open in the Events view, click the Column Group menu title and highlight the column group that you want edit.
2. Click  .
   The Edit Folder dialog is displayed.



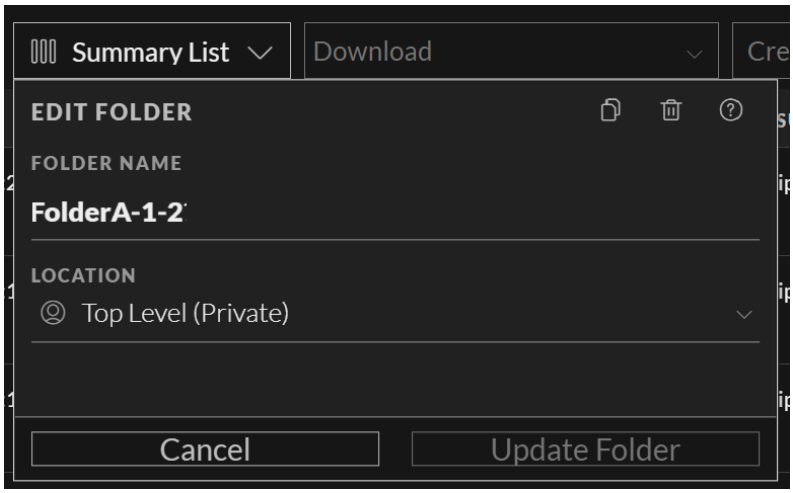
3. In the **Folder Name** field, type a unique name for the column group folder.
4. Select the location of the folder to be edited.
5. Click **Update Folder**.

Copy Column Group Folder

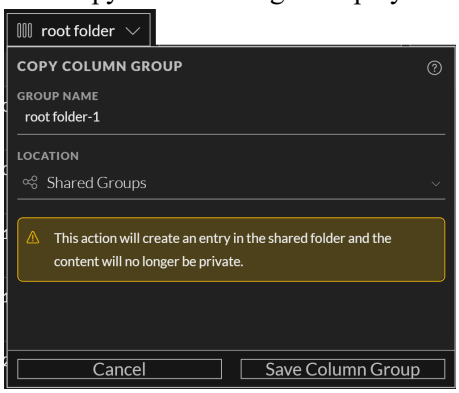
You can copy column group folders from private to shared, private to private, shared to shared and shared to private groups. When you copy a folder the content inside it gets copied. When you copy a private folder into a shared folder, the folder and its content no longer remain private.

1. With the Filter Events panel open in the Events view, click the Column Group menu title. The menu drops down to display a list of column groups and folders.
2. Select a folder you want to copy.

3. Click edit and the click copy .



The Copy Folder dialog is displayed.



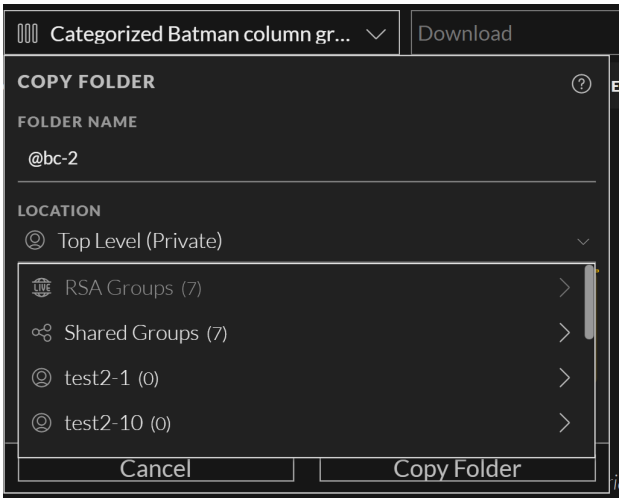
4. In the **Folder Name** field, type a unique name for the new meta group folder.
5. Select the location of the folder to be edited.
6. Click **Copy Folder**.

Copy Group Folder Deployed from Live

You can copy column group folder deployed from Live located under RSA Groups category to any other location like Shared groups or to a private folder.

1. With the Filter Events panel open in the Events view, click Column Group menu title. The menu drops down to display a list of column groups and folders.
2. Click on a Live Column Group folder you want to copy.

3. Click 
   The Copy Folder dialog is displayed.

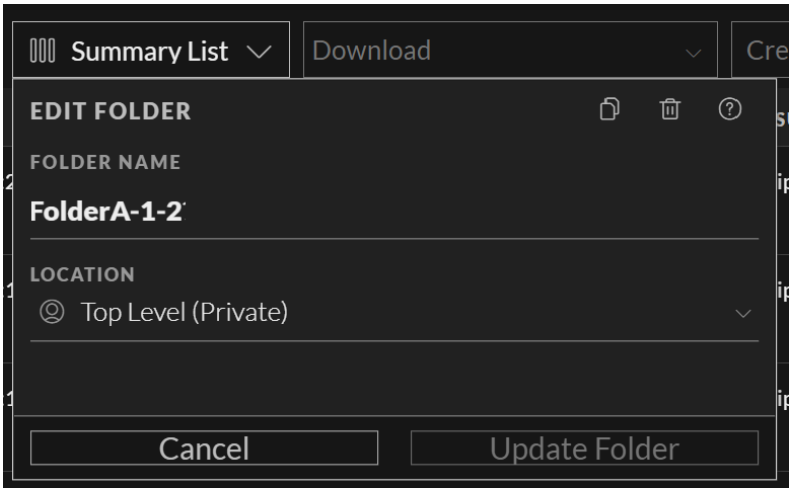


4. Select the location of the folder to be copied.
5. Click **Copy Folder**.
   The folder is created with the original name of the folder and appended with the 'copy' in the end.

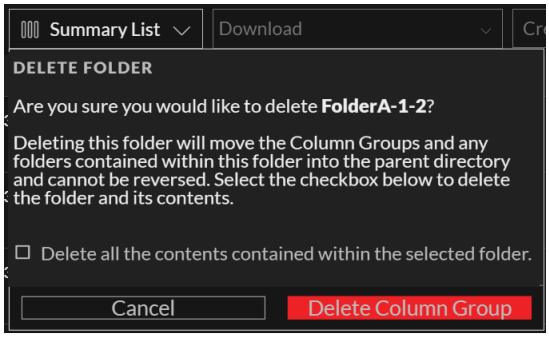
Delete Column Group Folder

If you don't want to retain a folder you can delete it. However, once the folder is deleted it cannot be retrieved.

1. With the Filter Events panel open in the Events view, click the Column Group menu title. The menu drops down to display a list of column groups and folders.
2. Select a folder to be deleted.
3. Click edit  .
   The Edit Folder dialog is displayed.



4. Click delete .
   A warning message is displayed to confirm the action.



5. (Optional) Select the checkbox, if you want to delete the folder along with all the contents inside the selected folder.
   If you do not select the checkbox, then the content will be moved to the parent folder after the required folder is deleted.
6. Click **OK** to delete.

Select a Column Group and Columns (11.3 and Earlier Event Analysis View)

In the 11.3 and earlier Event Analysis view, you can select a column group to apply to the Events list. These are the built-in groups and custom column groups created in the Legacy Events view.

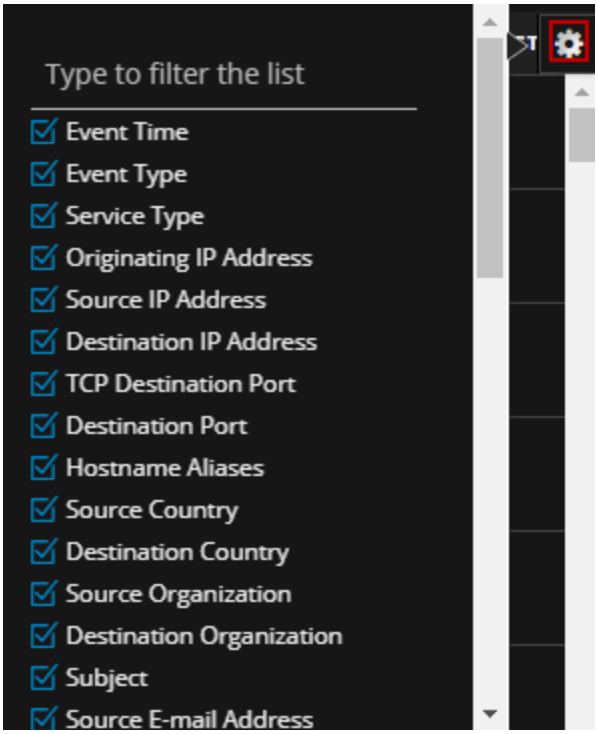
To select a column group:

Begin in the Column Group menu and do one of the following:

1. Select a column group (for example, **Summary List**).
2. To filter the list of column groups, start typing the name of the column group. Type one character and a list of column groups that contain that character is displayed; as you continue to type the list is filtered to match. When you see the column group that you want, click to select it. To clear the filter text, click **X** or delete the typed text.
   The Events panel displays data in the columns that belong to the selected column group.

To select columns to display:

1. With the Events list open and a column group selected, click ⚙ to display the column selector.



2. Select the meta keys or enter the name of a meta key that you want to display in additional columns.
3. If you do not want to see a meta key displayed in a column, deselect the meta key.
   The data is redisplayed using the selected columns.

Work with Column Groups in the Legacy Events View

This section includes procedures for working in the 11.4 Legacy Events view (and the 11.3 Events view). Three different forms of the events list with hard-coded columns are built in and labeled as follows: Detail View, List View, Log View. You can remove columns, rearrange the order, and change the width of a column. In addition the built-in or custom column groups are available; these give you more flexibility in choosing columns.

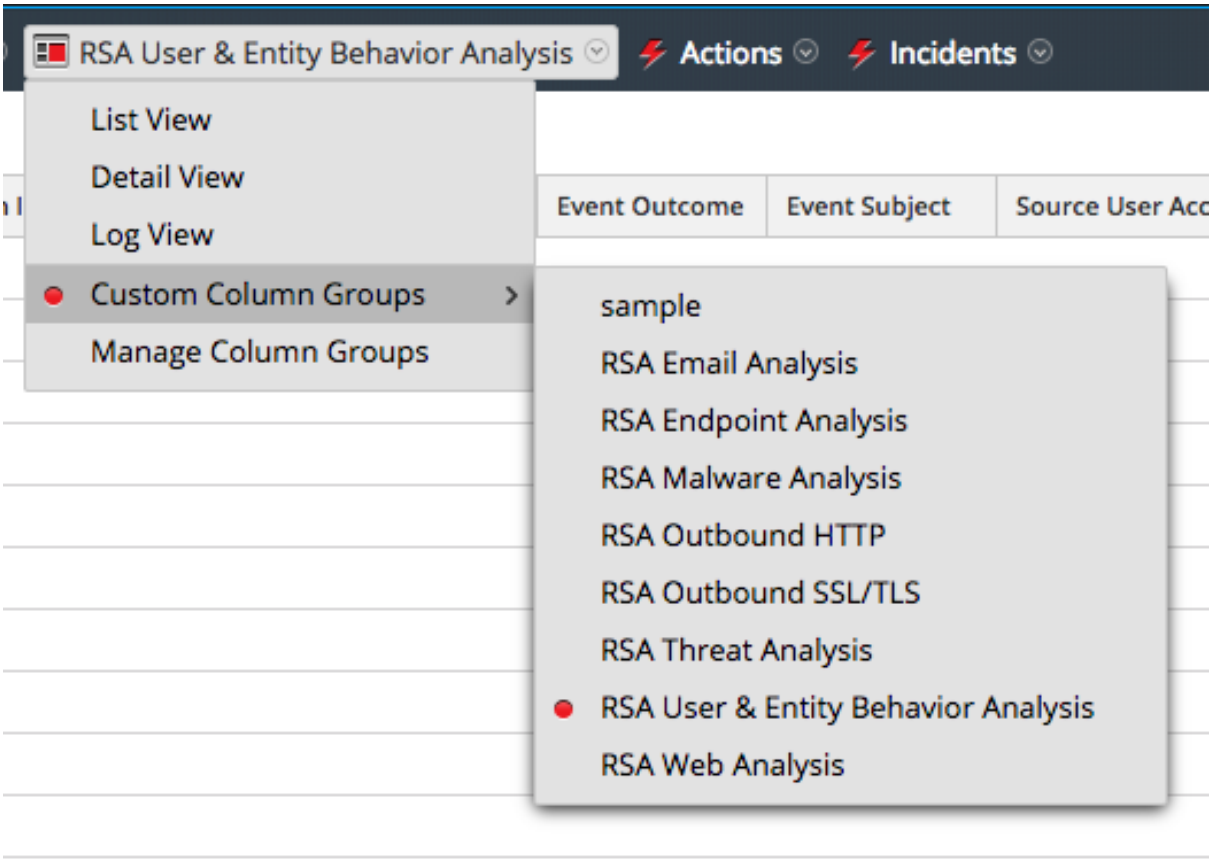
Column groups are shared globally, per service, across Investigate. Any changes you make to custom column groups are applied globally, affecting all analysts using the service. If you delete a column group, the column group is no longer available to anyone who is investigating the service.

Select a Column Group

Note: Investigate profiles can include custom column groups. If a custom column group is used in a profile and you are viewing events in the Legacy Events view using a custom column group, you cannot change the view type (Detail, List, or Log).

To select a column group:

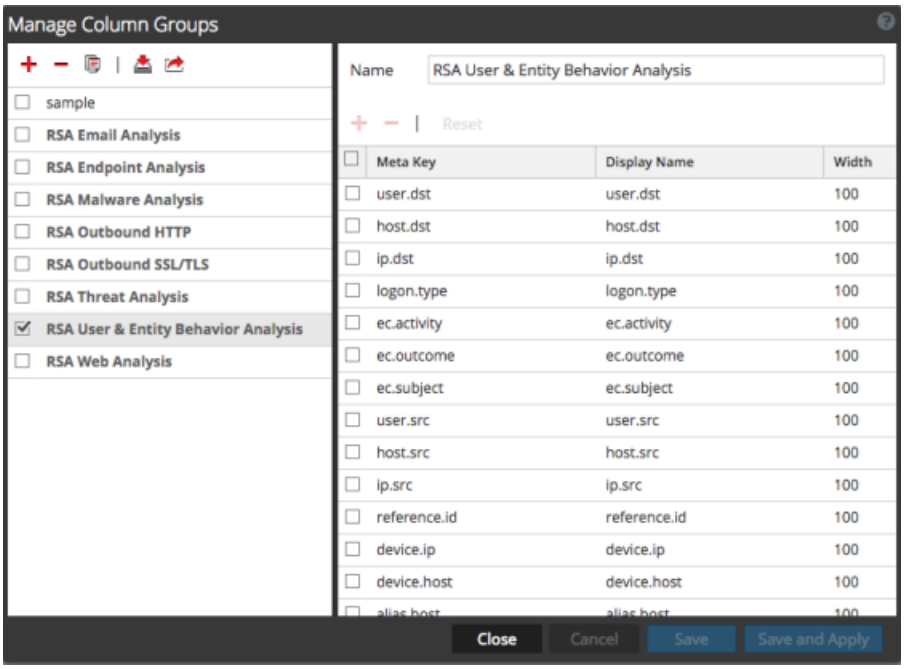
1. With the Legacy Events view open, select **Custom Column Groups** in the **View** drop-down menu. The menu label reflects the selected option: Detail View, List View, Log View, or the currently selected column group.



2. Select one of the column groups from the submenu.
   The Legacy Events view is refreshed to reflect the custom column group.

Create a Custom Column Group in the Legacy Events View

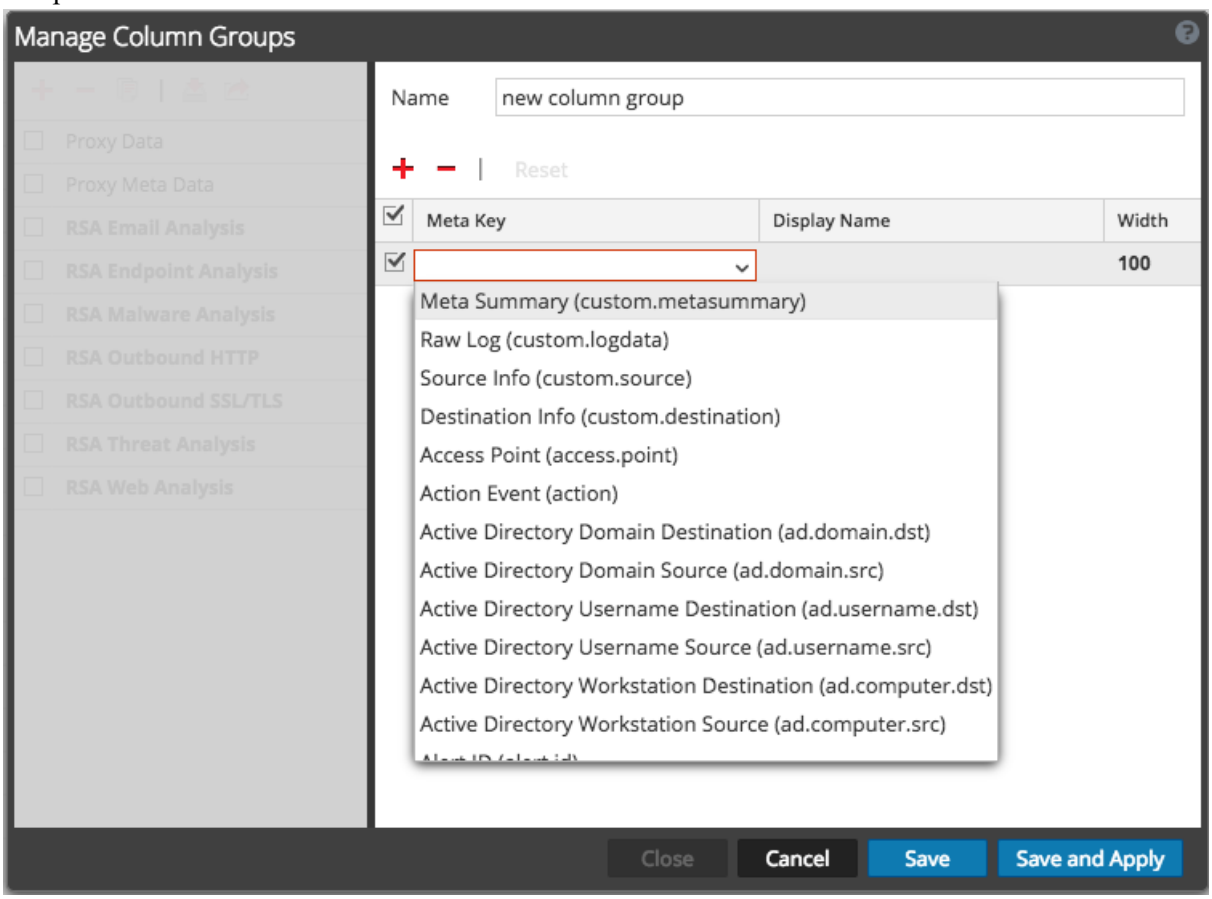
1. Go to **Investigate > Legacy Events**.
2. Select **Manage Column Groups** in the **View** drop-down menu. The View option is named for the current value, for example, Detail View, List View, Log View, or the currently selected column group.
   The Manage Column Groups dialog is displayed.



3. To add a new column group in the column group panel, click + and type the name of the new group in the resulting field.
   The column definition panel opens on the right with the group name filled in. You can edit the group name.

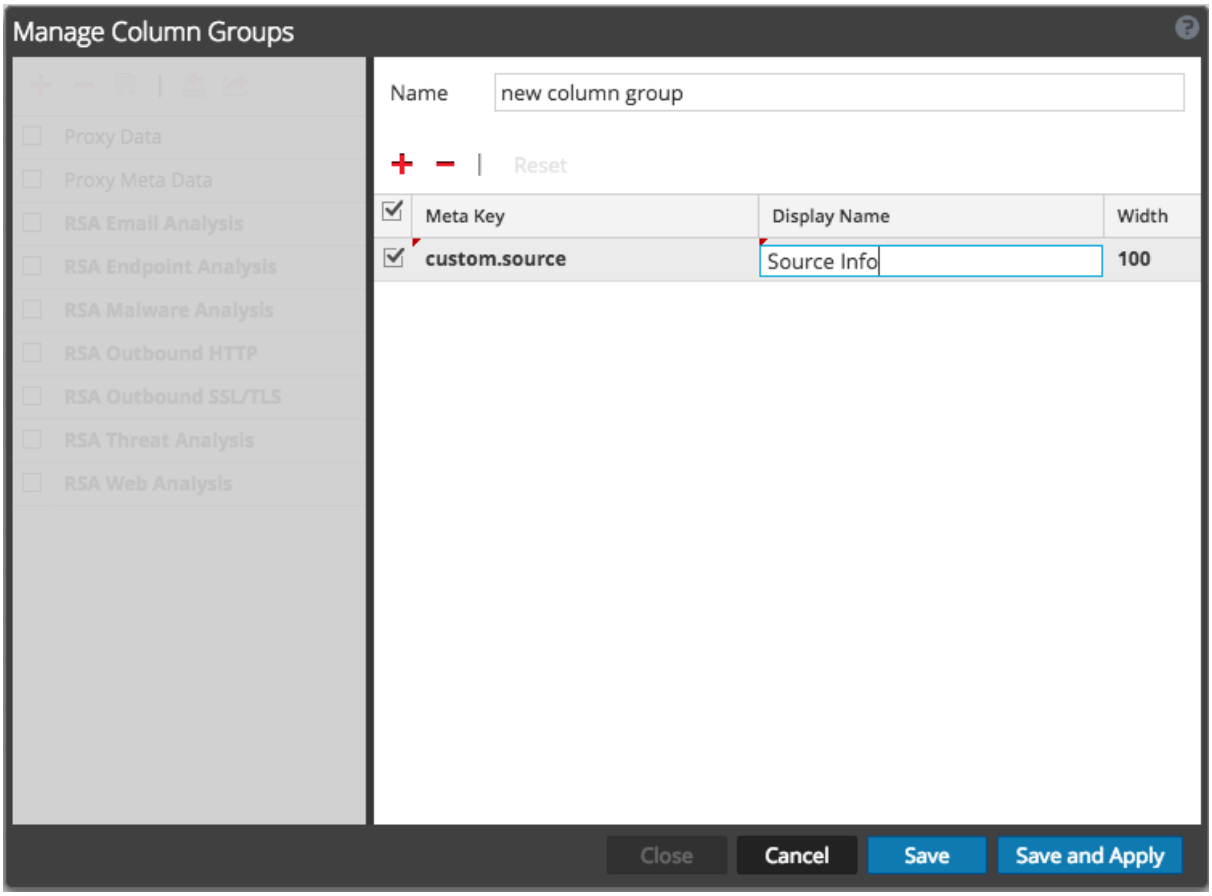
4. To add a column to the group, click +, and click in the empty **Meta Key** field to display the **Meta Key** drop-down list. Select a meta key field from the list, and repeat this step until the column set is

complete.



5. (Optional) To delete a meta key from the column group, click −.
6. (Optional) To rearrange the sequence in which the columns appear in the Events list, drag meta keys to the desired position.

7. (Optional) To set the default width for a column, click in the corresponding value in the **Width** column, and type a new column width.

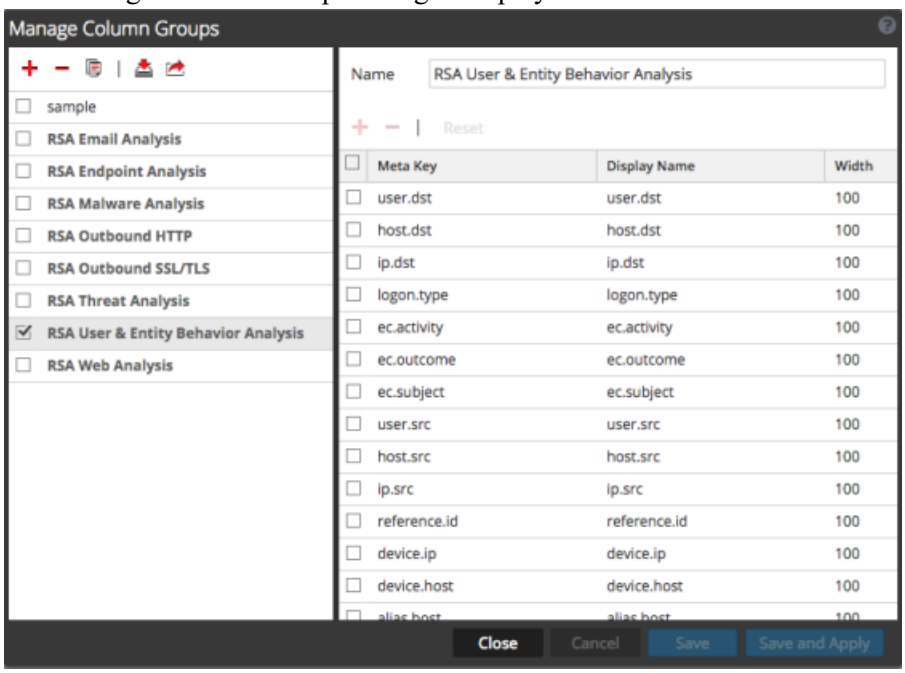


8. (Optional) To revert to the previous settings for the column group, and undo all of your changes, click **Cancel**.
9. When ready to save, do one of the following:
   a. To save the edited column group and refresh the Legacy Events view with the column group settings, click **Save and Apply**.
   b. To save the edited column group without refreshing the Legacy Events view, click **Save**.

Delete a Column Group (Legacy Events View)

1. Go to **Investigate > Legacy Events**.
2. Select **Manage Column Groups** in the **View** drop-down menu. The View option is named for the current value, for example, Detail View, List View, Log View, or the currently selected column group.

The Manage Column Groups dialog is displayed.

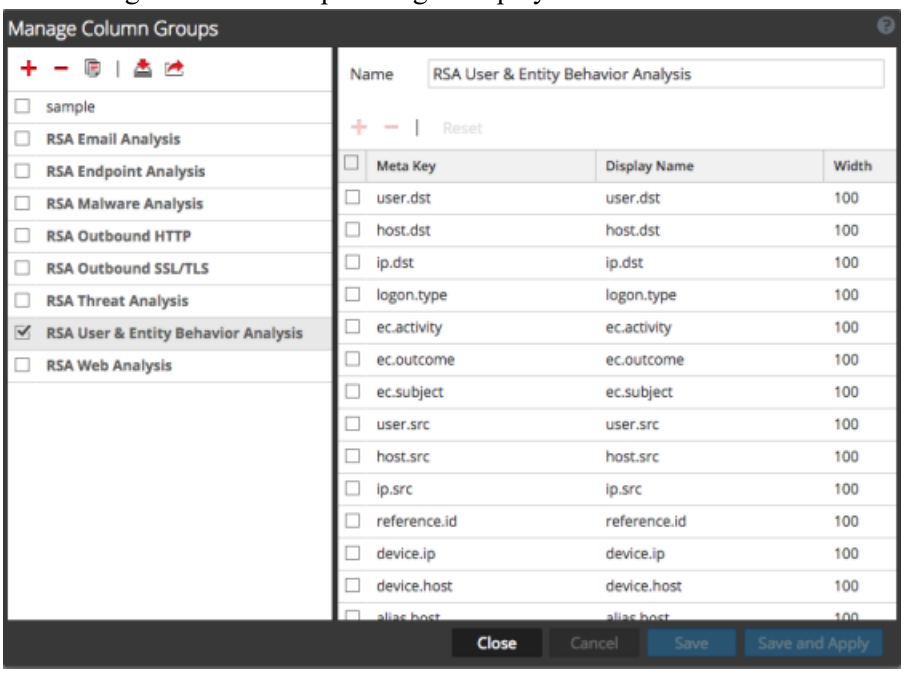


3. To delete a custom column group in the column group panel, select one or more custom column groups and click − in the toolbar.
   A confirmation request is displayed.
4. Do one of the following:
   a. To delete the column group and refresh the Legacy Events view, click **Yes**.
   b. If you decided not to delete the column group , click **No**.
      The selected column groups are deleted and no longer appear anywhere for this service in Investigate.

Edit a Column Group (Events View)

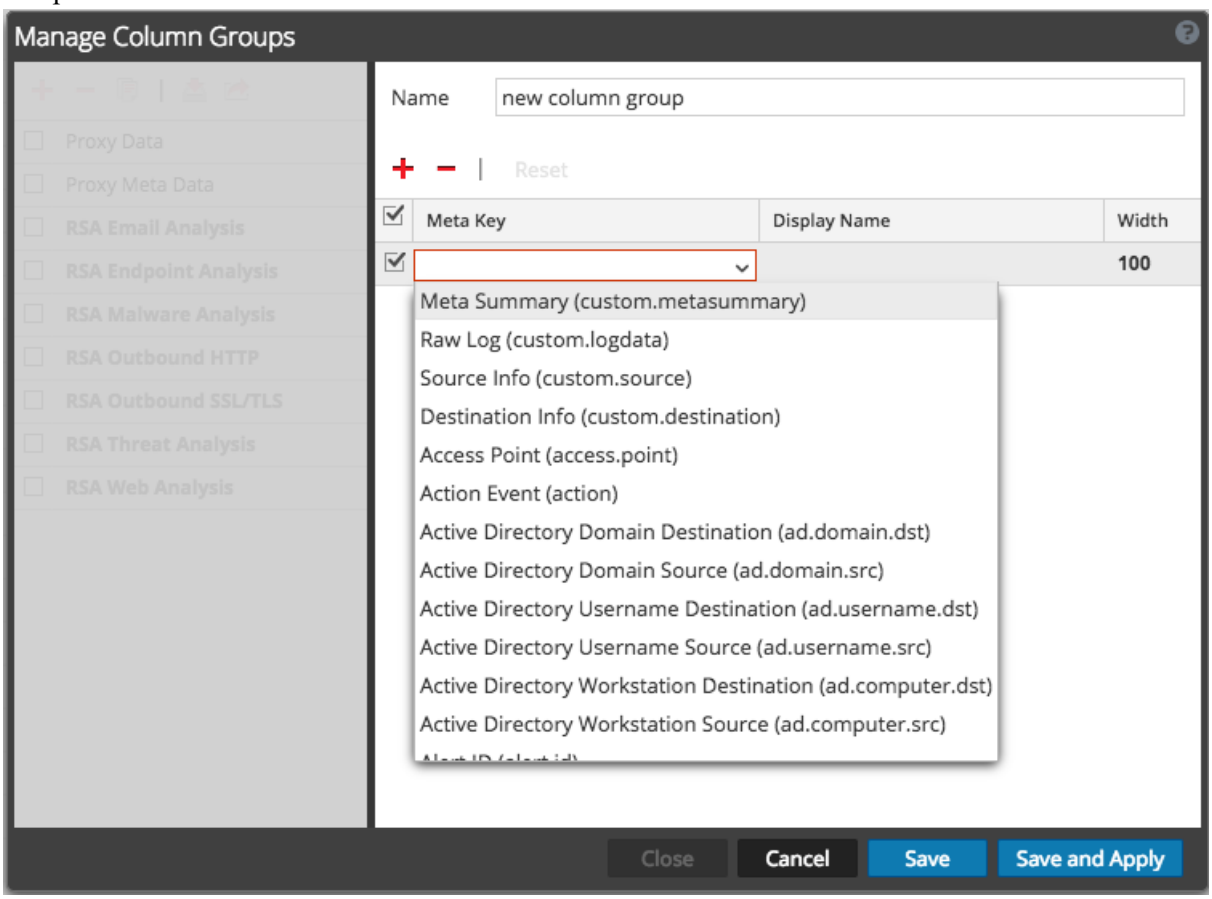
1. Go to **Investigate > Legacy Events**.
2. Select **Manage Column Groups** in the **View** drop-down menu. The View option is named for the current value, for example, Detail View, List View, Log View, or the currently selected column group.

The Manage Column Groups dialog is displayed.



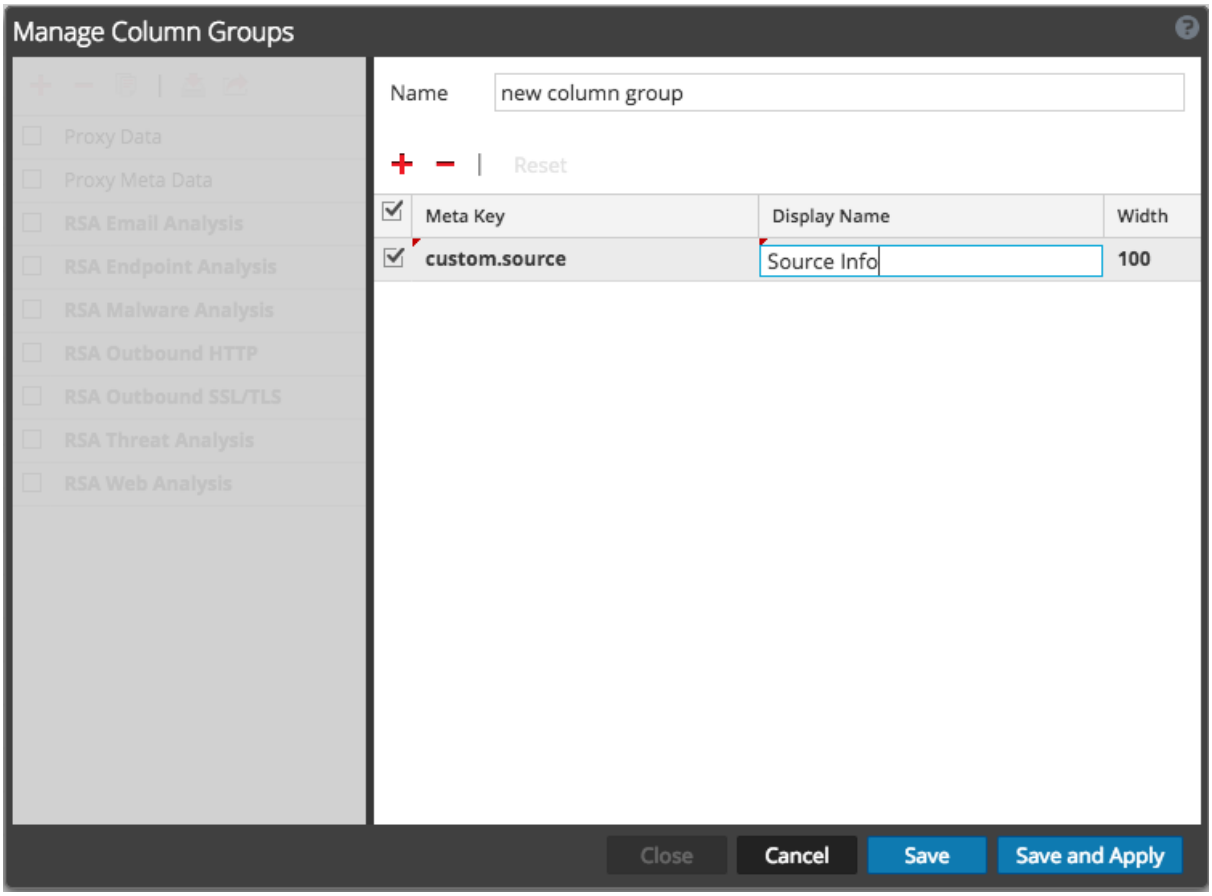
3. Do one of following:
   a. To edit a custom column group in the column group panel, select the checkbox before the name. The column definition panel opens on the right.
   b. To clone and edit a built-in column group or a custom column group, select the checkbox before the name and click the clone icon ( ).
   The column definition panel opens on the right.
4. (Conditional) If you are editing a clone of a group, type the new name of the group.
5. To add a column to the group, click , and click in the empty **Meta Key** field to display the **Meta Key** drop-down list. Select a meta key field from the list, and repeat this step until the column set is

complete.



6. (Optional) To delete a meta key from the column group, click −.
7. (Optional) To rearrange the sequence in which the columns appear in the Events list, drag meta keys to the desired position.

8. (Optional) To set the default width for a column, click in the corresponding value in the **Width** column, and type a new column width.



9. (Optional) To revert to the previous settings for the column group, and undo all of your changes, click **Cancel**.
10. When ready to save, do one of the following:
    a. To save the edited column group and refresh the Legacy Events view with the column group settings, click **Save and Apply**.
    b. To save the edited column group without refreshing the Legacy Events view, click **Save**.

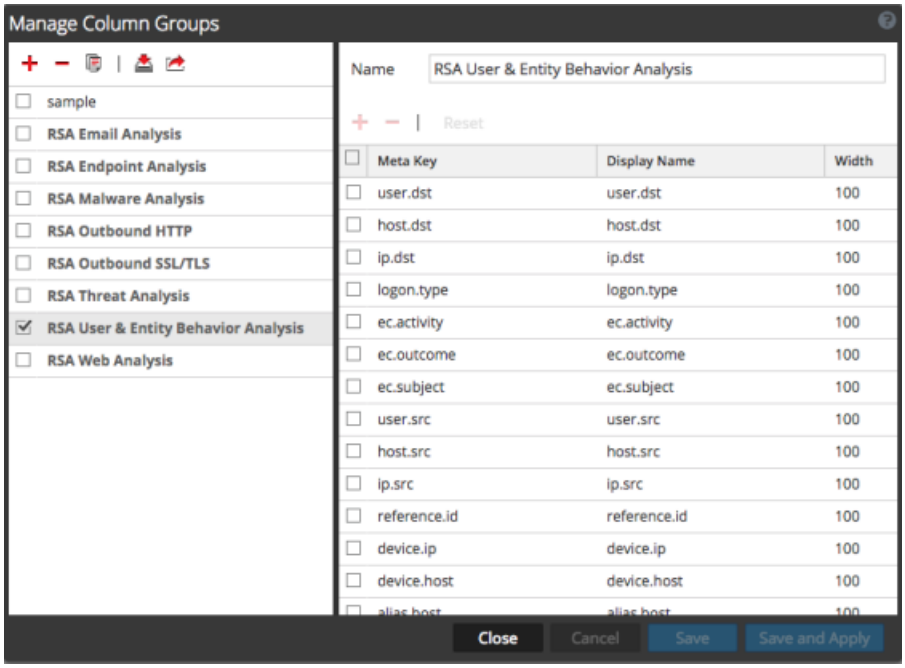
Import and Export a Column Group (Legacy Events View)

You can export custom column groups for use by other members of your team, and other analysts can import column groups if you give them a copy of the exported file.

To export a column group:

1. Go to **Investigate > Legacy Events**.
2. Select **Manage Column Groups** in the **View** drop-down menu. The View option is named for the current value, for example, Detail View, List View, Log View, or the currently selected column group. Each of these views is a differently formatted events list, and each column represents one

meta key.
The Manage Column Groups dialog is displayed.



3. To export a column group, select the checkbox before the name and click the Export option ( ). The column group is exported to your local file system as a .jsn file, for example, `CustomColumnGroupsExport.jsn`. If you export another group, the next file is named `CustomColumnGroupsExport-2.jsn` to differentiate.
4. To import a column group that you have available on your local file system, click the Import option ( ).
   The Import Column Groups dialog is displayed.
5. Browse your local drive to find the column group (jsn file), and click **Upload**.
   The column group is added to the list. If it has the same name as an existing column group, a message is displayed and the column group is not imported.

Use Query Profiles to Encapsulate Common Areas for Investigation

Query profiles offer a quick and easy way to define a meta group, column group, and a limiting filter (pre-query condition) that you can apply in the Navigate view, the Events view, and the Legacy Events view. The same query profiles are shared between all views, and they are available in the Springboard (Version 11.5) for use in panels. Private query profiles created in the Events view are only available in the Events view for the analyst who created them.

Each query profile specifies a meta group, column group, and sometimes includes a pre-query condition appropriate for the type of investigation.

In a query profile:

- The meta group defines the meta keys that are queried (see Use Meta Groups to Focus on Relevant Meta Keys).
- The column group defines which meta keys from the meta group are displayed as columns in the Events list. (see Use Columns and Column Groups in the Events List).
- When the query profile is in effect, the optional pre-query conditions add a limiting filter in the query bar. You can edit or delete the limiting filter and then create additional filters for your query (see Filter Results in the Events View).

Built-In Query Profiles

You cannot edit or delete built-in profiles, but you can copy an existing profile and edit the copy in the Navigate view, the Legacy Events view, or the Events view. In the Navigate view, the built-in profile names begin with the RSA prefix and are grouped under Default Profiles. The Events view does not support grouping of query profiles. This figure is an example of a built-in query profile as listed in the Query Profiles menu.



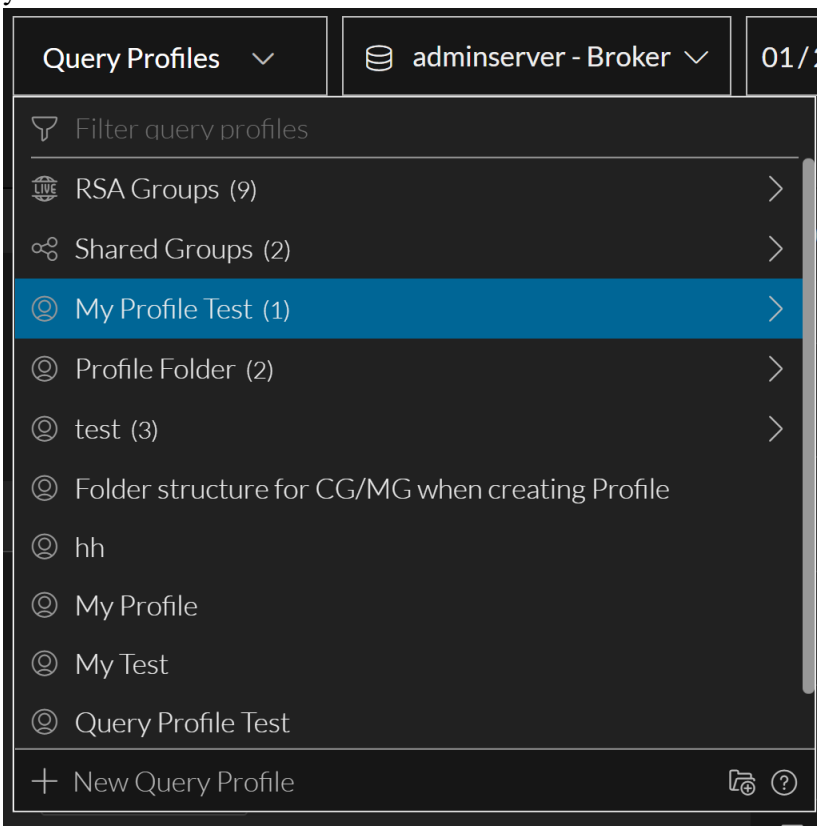
NetWitness Platform has these built-in profiles:

- RSA Email Analysis
- RSA Endpoint Analysis
- RSA File Analysis
- RSA Threat Analysis
- RSA User & Entity Behavior Analysis
- RSA Web Analysis

Built-in query profiles make it easy for you to query a specific area of interest; for example, selecting the built-in RSA Email Analysis query profile automatically specifies the meta group, and column group, and pre-query conditions that are most useful for investigating email activity. As you become familiar with the meta keys, you can create your own custom query profiles.

Live Query Profiles

In 11.6 and later, NetWitness supports deploying the investigate content from live and are marked by the live symbol () under the query profiles group drop down. The query profiles are categorized as RSA Groups (RSA Live content and RSA OOTB Groups), and Shared Groups. The groups are displayed as non-editable folders and sub-folders except for Shared Groups that can be edited. All private content is displayed outside these groups. For example, the below image shows private content below the Shared Groups folder. The number inside () depicts the number of contents inside a folder and > symbol helps you to drill down inside the folder.



Custom Query Profiles

Custom query profiles are shared globally within your organization in Version 11.4. In Version 11.5 and later, you can create shared query profiles as before, and can also create private query profiles. If you edit a shared custom query profile, your changes are applied globally. If you delete a shared custom query profile, the profile is deleted and no longer available for all analysts.

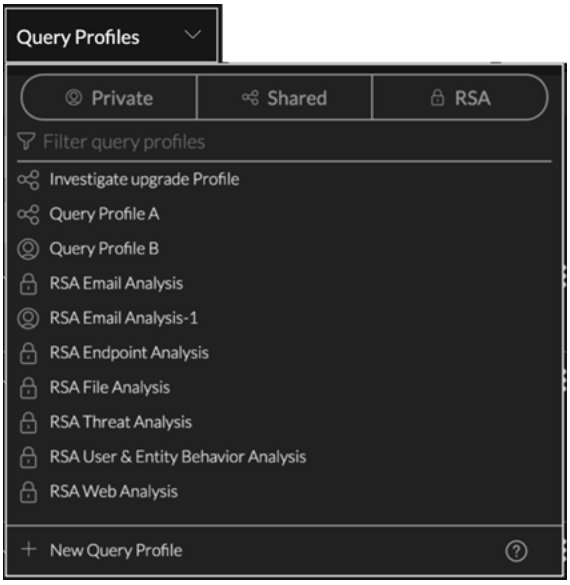
Note: If a Springboard panel is using a query profile as a filter, the profile can be edited, but cannot be deleted in the Events view. However, nothing prevents deletion of the profile in the Navigate view or the Legacy Events view. In this case, Springboard panels that use the deleted query profile as a filter continue to work, but the filter is removed and unexpected results may be displayed in the panel. Refer to "Managing the Springboard" in the *NetWitness Platform Getting Started Guide* for details.

When you create a query profile in Version 11.5, you can choose to share it or you can keep it private (default); you cannot change a shared profile to private or a private profile to shared. Private query profiles are not visible or usable in the Navigate view, the Legacy Events view, or the Springboard. Icons identify the profile type in the Query Profile menu. These are examples of a shared and a private custom query profile as listed in the Query Profile menu, with the edit icon displayed at the end of the row.

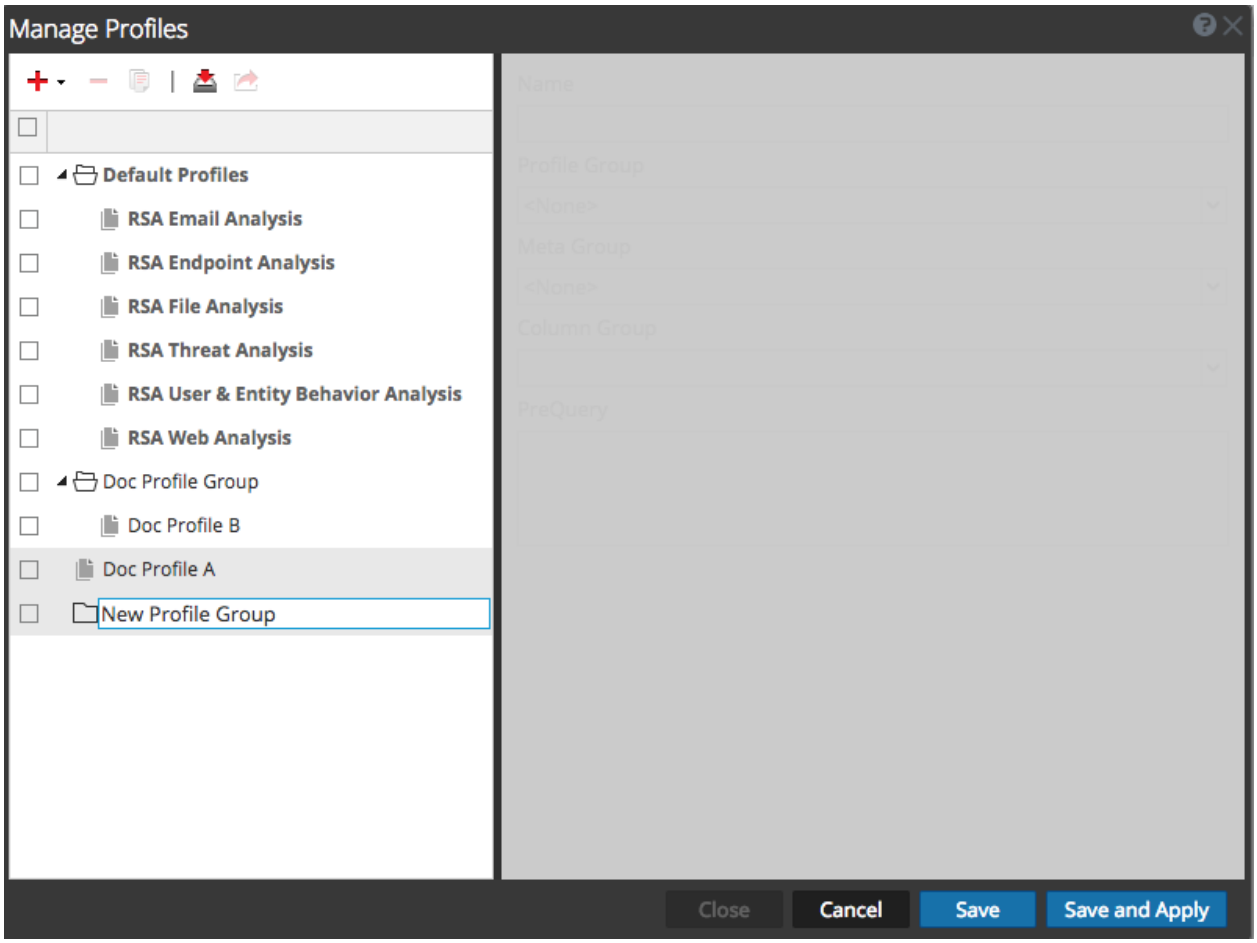


Dialogs for Managing Query Profiles

The profiles are listed in alphabetical order in the Query Profiles menu in a way that makes built-in profiles distinguishable from custom profiles that you imported or created. While the functionality for managing query profiles is similar in the Navigate view, the Legacy Events view, and the Events view, the dialogs are different. The following figure illustrates the Query Profiles menu in the Version 11.5 Events view. This menu lists the same profiles that are available in the Navigate view and the Legacy Events view. You can create, copy, edit, delete, and apply profiles. You can filter the list of profiles using the filtering buttons (Private, Shared, and RSA) at the top to display any combination of private, shared, and built-in query profiles.



This is an example of the Manage Profiles dialog in the Navigate and Legacy Events views.



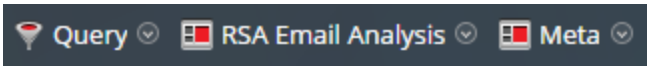
Note: Query profiles are available in the Navigate view, the Legacy Events view, and the Events view; in Version 11.4.1 and earlier, they are shared globally across users. If one user modifies or deletes a custom query profile it has an effect on what is available to the other users. In the Events view, use the Query Profiles menu to work with profiles. In the **Navigate** or **Legacy Events** view toolbar, select **Profile > Manage Profiles** to open the Manage Profiles dialog. In Version 11.5, custom profiles can be shared globally, but private custom profiles created in the Events view are not available in the Navigate view or the Legacy Events view.

From the Query Profiles menu (11.4 and later Events view):

- You can apply a query profile and use options in the menu to create (Create Query Profile dialog), copy, edit, and delete (Query Profile Details dialog) custom query profiles.
- Selecting a profile applies the meta group, column group, and pre-query condition, and these are visible in the Meta Group menu title, Column Group menu title, and the query bar.
- In Version 11.4, the Events view does not use meta groups or profile groups defined in other views. Version 11.5 allows you to use meta groups and to create private custom query profiles, in addition to the previously available shared custom query profiles.
- If a query profile created in the Legacy Events view uses the Log View, Detail View, or List View instead of a column group, the same profile in the Events view uses the Summary List column group.

From the Manage Profiles dialog (Navigate view and Legacy Events view):

- You can configure, add, delete, import, and export profiles and profile groups.
- You can organize your custom query profiles in profile groups (Version 11.2 and later). When upgrading to Version 11.4 from an earlier version, only profile groups that contain profiles are imported. The built-in query profiles are in the Default Profiles group, which cannot be edited. Analysts can create new query profile groups, which anyone can use.
- After creating profiles, you can edit a profile group to add profiles, remove profiles, or move profiles from one group to another. When you create a profile, it is not added to any profile group by default.
- Selecting a profile applies the meta group, column group, and pre-query condition, and the label of the Profile menu is replaced with the query profile name. The following figure illustrates the RSA Email Analysis query profile selected in the Navigate view or Legacy Events view.

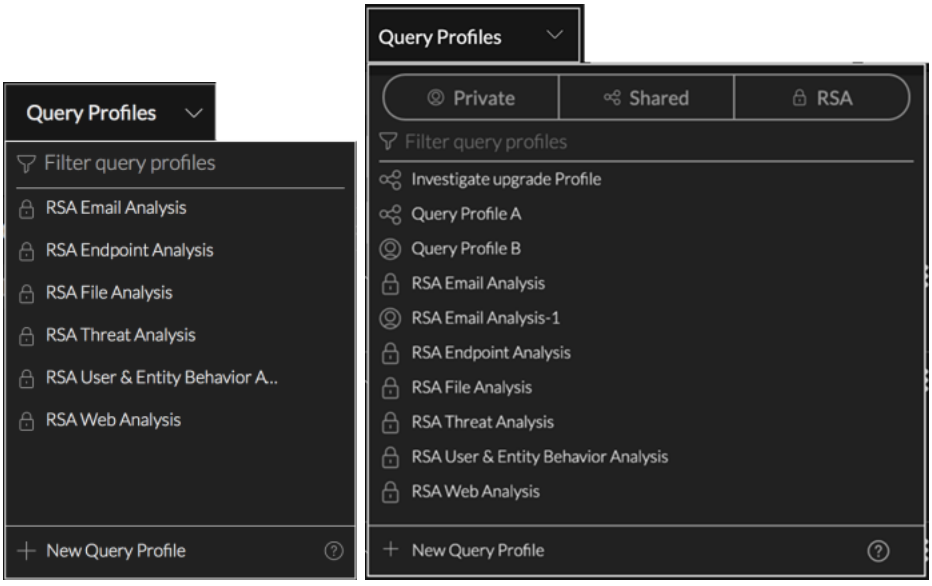


View Query Profile Details (Events View)

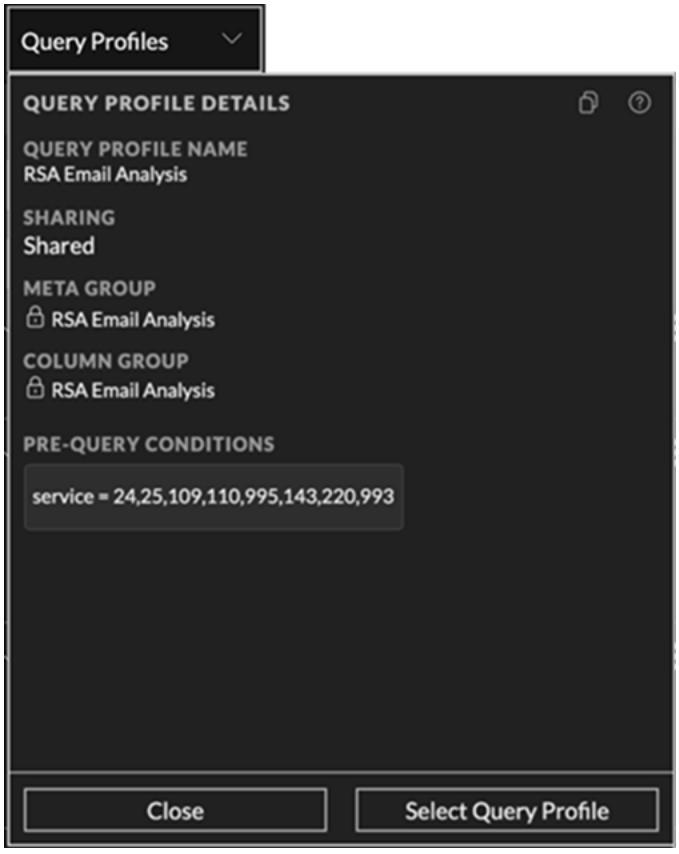
If you want to know which meta groups, column groups, and limiting filters (called pre-query conditions) define a query profile, you can view the details of the profile.

To view the details:

1. Go to **Investigate > Events** and click **Query Profiles** in the query bar.
   The Query Profiles menu opens with a list of available profiles. The Version 11.5 menu displays a list a list of built-in query profiles (RSA), shared custom profiles, and your private custom profiles with visibility options and a filter field make it easier to find a particular query profile. The figure on the left shows the Version 11.4 menu. The figure on the right shows the Version 11.5 menu initially with all types of profiles visible: Private, Shared, and RSA.



2. Hover over a query profile in the list and click the information icon ( ) to see the meta group, column group, and pre-query conditions configured for the profile.
   This figure shows the details for the RSA Email Analysis profile, one of the built-in profiles. In Version 11.5.1, an icon identifies the type of meta group and column group (shared, private, or RSA).



3. Do one of the following:

   a. To close the dialog, click **Close**.

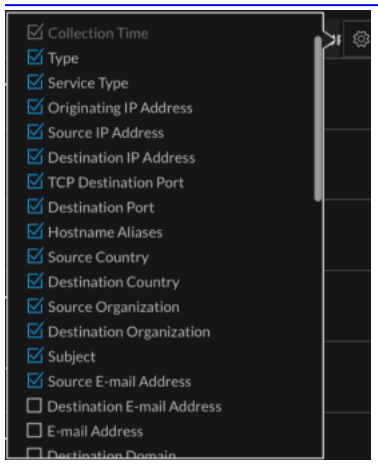
   b. If you want to apply the profile, click **Select Query Profile**.
      The dialog closes. The Events list is updated to reflect the selected query profile. If the profile uses a different column group, the query is re-executed with the pre-query conditions and column group for the selected profile. If only the pre-query conditions are different, existing filters in the query bar are removed and the pre-query conditions (for example, this filter: `service=24,25,109,110,995,143,220,993`) is added in the query bar, but the query is not submitted. The first 15 columns in the associated column group are used in the Events list.

      i. (Optional) Create additional filters in the query bar before executing the query (see Filter Results in the Events View).



      ii. (Optional) If you want to select different columns from the associated column group before executing the query, click above the Events list on the right.
      The Column Selection list is displayed and you can choose up to 40 columns to display (see

Use Columns and Column Groups in the Events List.



Apply a Query Profile (Events View)

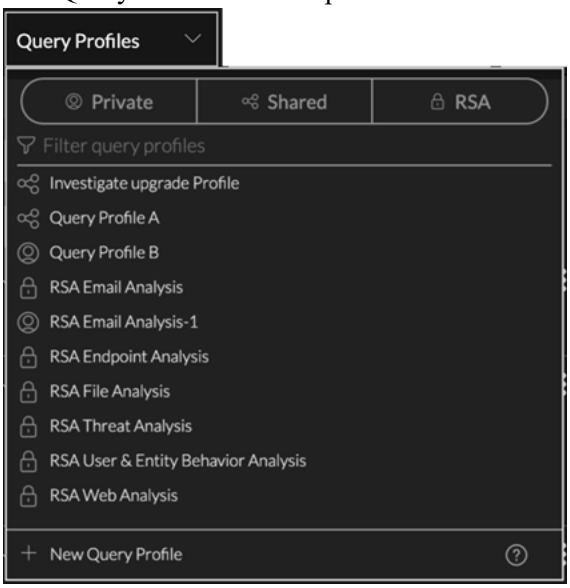
When a query profile is applied, there is no indication of it in the Query Profile menu, but you can see if a column group or meta group is in effect. If pre-query conditions are applied, the filters are visible at the beginning of the query bar as shown in this figure:



Note: If you do not see enough results or the right results in the Events view, an applied profile may be limiting results with pre-query conditions.

To apply a query profile:

1. Go to **Investigate > Events** and click **Query Profiles** in the query bar.
   The Query Profiles menu opens with a list of available profiles.



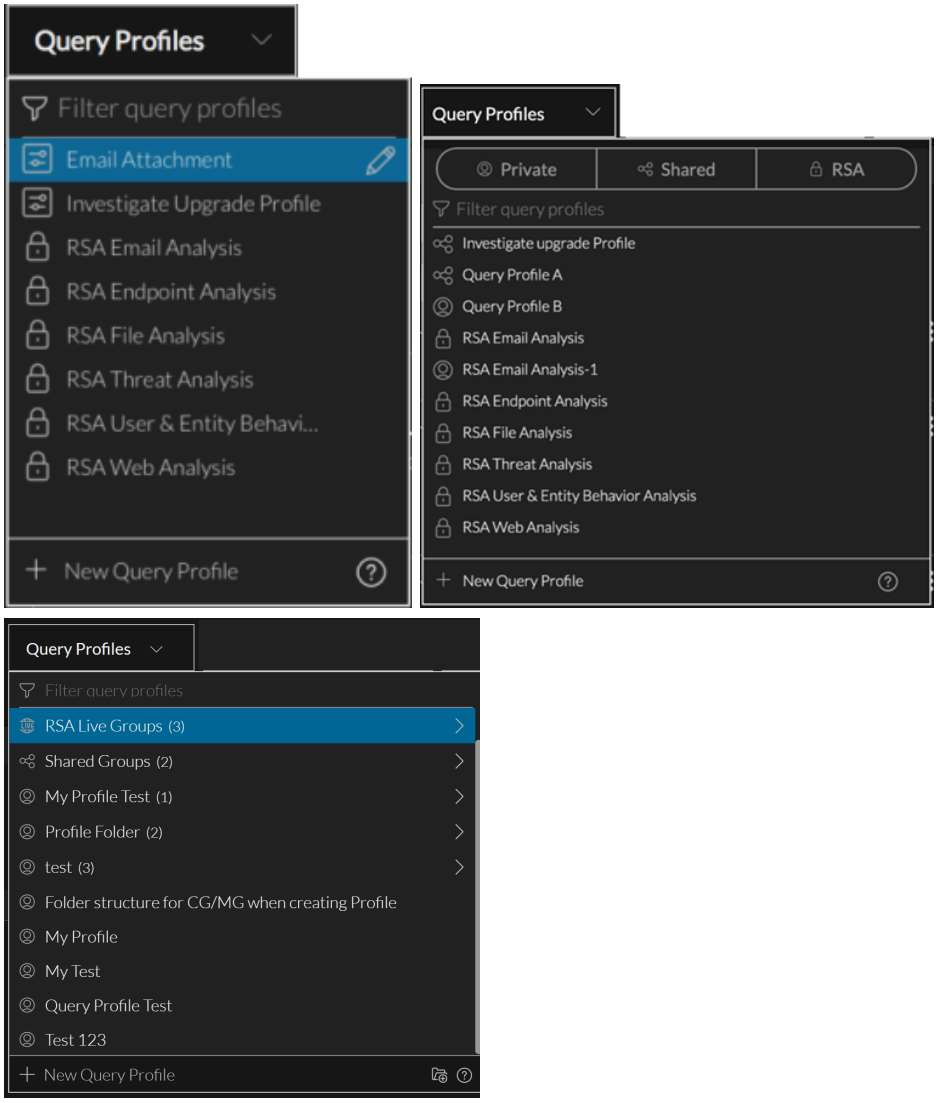
2. Use the Down and Up arrow keys or the mouse to highlight a profile.

3. Click the highlighted profile.
   The query profile settings are applied immediately. The Events list is updated to reflect the selected profile. If the profile uses a different column group the query is re-executed with the pre-query conditions and column group for the selected profile. If only the pre-query conditoins are different, existing filters in the query bar are removed and the pre-query conditions are added in the query bar. The button becomes active so that you can resubmit the query with the new pre-query conditions. You can add more filters as usual before or after resubmitting the query.

Create or Edit a Custom Query Profile (Events View)

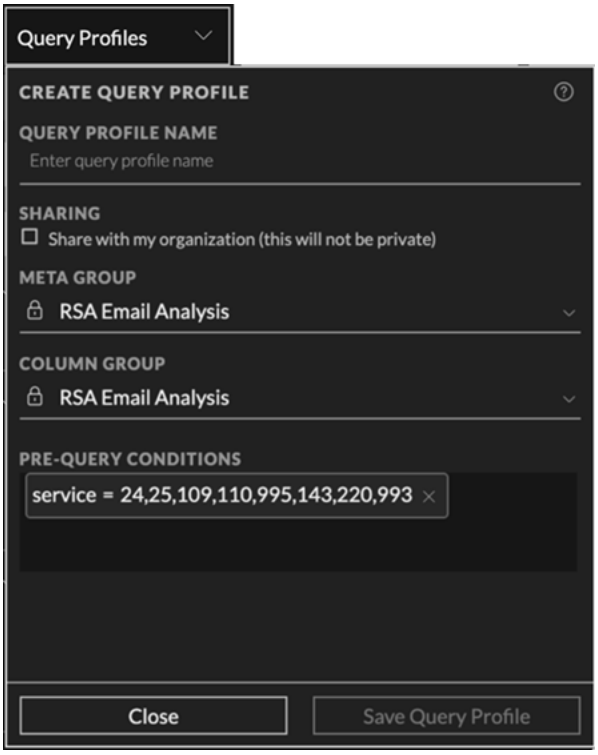
To create or edit a custom query profile:

1. Go to **Investigate > Events** and click **Query Profiles** in the query bar.
   The Query Profiles menu opens with a list of available profiles. Version 11.6 (figure on the right) and version 11.5 (figure in the center) supports both private and shared custom profiles; Version 11.4 (figure on the left) supports only shared custom profiles.

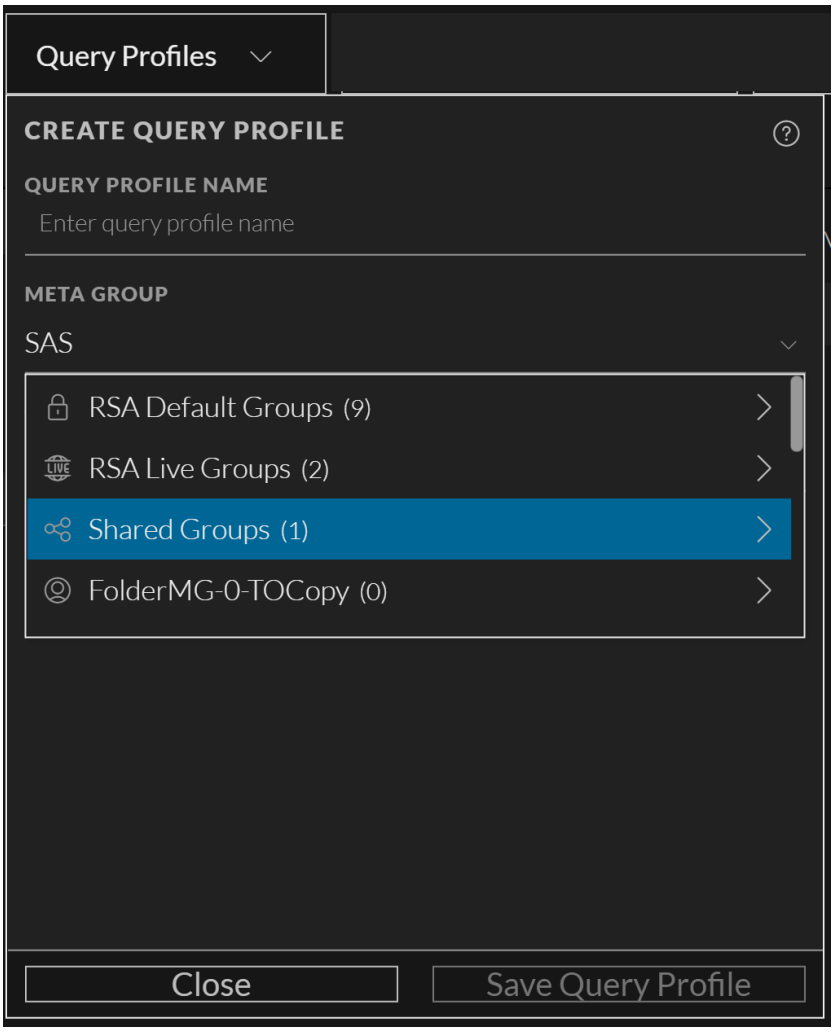


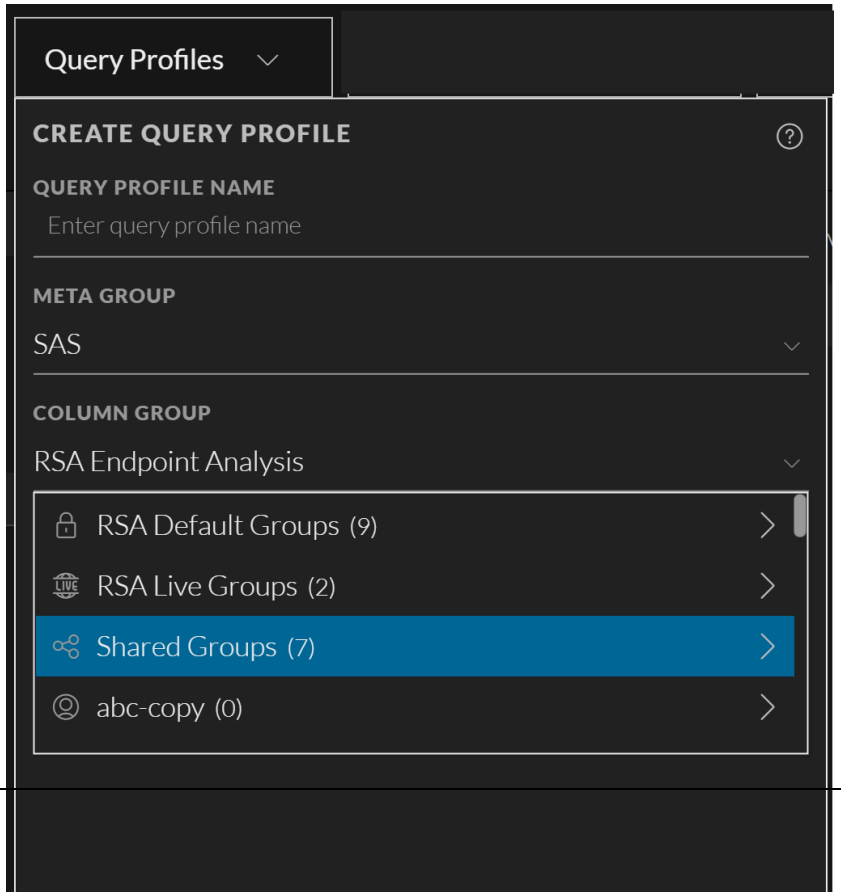
2. Do one of the following:

   a. To create a new query profile, click + **New Query Profile**.
   The Create Query Profile dialog is displayed. The Create Query dialog shows a new empty profile that includes the currently selected meta group, column group, and filter that you have currently typed in the Query bar as a pre-query condition.

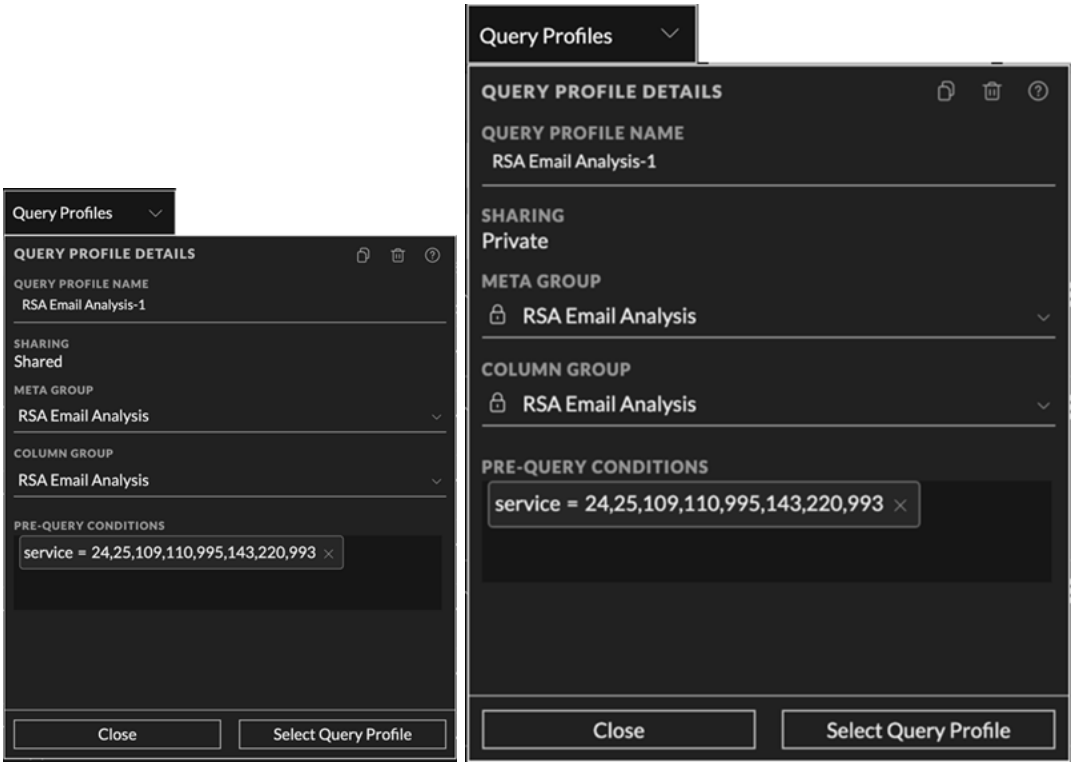


   In version 11.6, the meta groups and column groups are displayed in a drop-down in a folder structure.


Query Profiles
CREATE QUERY PROFILE
QUERY PROFILE NAME
Enter query profile name
META GROUP
SAS
RSA Default Groups (9)
RSA Live Groups (2)
Shared Groups (1)
FolderMG-0-TOCopy (0)
Close
Save Query Profile


Query Profiles
CREATE QUERY PROFILE
QUERY PROFILE NAME
Enter query profile name
META GROUP
SAS
COLUMN GROUP
RSA Endpoint Analysis
RSA Default Groups (9)
RSA Live Groups (2)
Shared Groups (7)
abc-copy (0)

b. To edit an existing query profile, highlight a custom query profile in the menu, and click the edit ( ) icon.
The Query Profile Details dialog is displayed. The Version 11.5.1 dialog (on the right) identifies the type of meta group and column group as shared, private, or RSA.



3. In the **Profile Name** field, type a unique profile name that has no more than 80 characters.
In the Create Query dialog, the Save Query Profile button is activated. In the Query Profile Details dialog, the Select Query Profile button is relabeled as Update Query Profile.
4. (Version 11.5 and later), do one of the following
   a. If you want to share the new query profile with your organization, set the **Share with my organization** option. You cannot change a query profile from shared to private after it is created.
   b. If you want to create a private query profile that only you can see and manage, leave the **Share with my organization** checkbox empty. You cannot change a query profile from private to shared after it is created.
5. (Version 11.5 and later) Select a meta group from the **Meta Group** drop-down list. If a shared group and a private group have the same name, the private group is listed before the shared group. In Version 11.5.1, an icon before the group name distinguishes private from shared.
6. Select a column group from the **Column Group** drop-down list. In Version 11.5, there can be shared or private groups and they can have the same name. In this case, the private group is listed before the shared group. In Version 11.5.1, an icon in front of the group name distinguishes private from shared.
7. In the **Pre-Query Conditions** field, check the default filters from the query bar and add or remove filters if you wish.

8. Click **Save Query Profile** or **Update Query Profile**.
   The new profile is saved or the edited profile is updated with your changes.
9. To close the dialog, click **Close**.

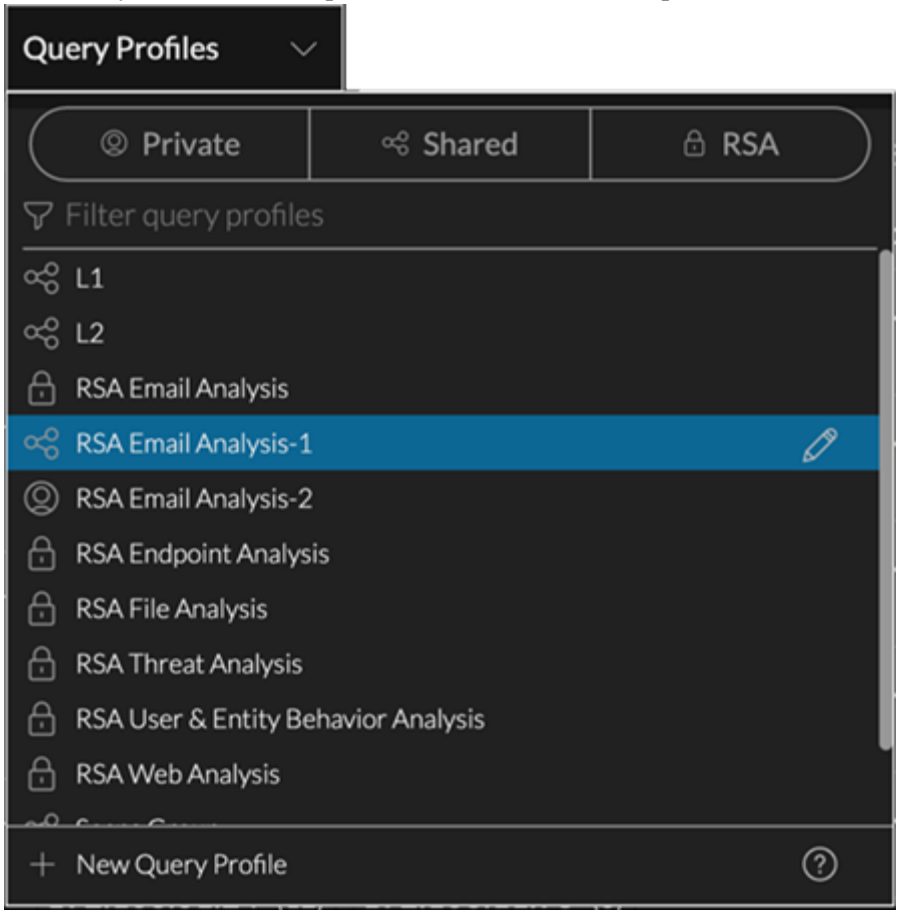
Delete a Custom Query Profile (Events View)

Built-in query profiles are read only, and cannot be deleted, but you can delete any custom query profile. A confirmation message allows you to confirm or cancel the deletion. When you delete a shared query profile, the effect is global and the profile is no longer available to any analyst.

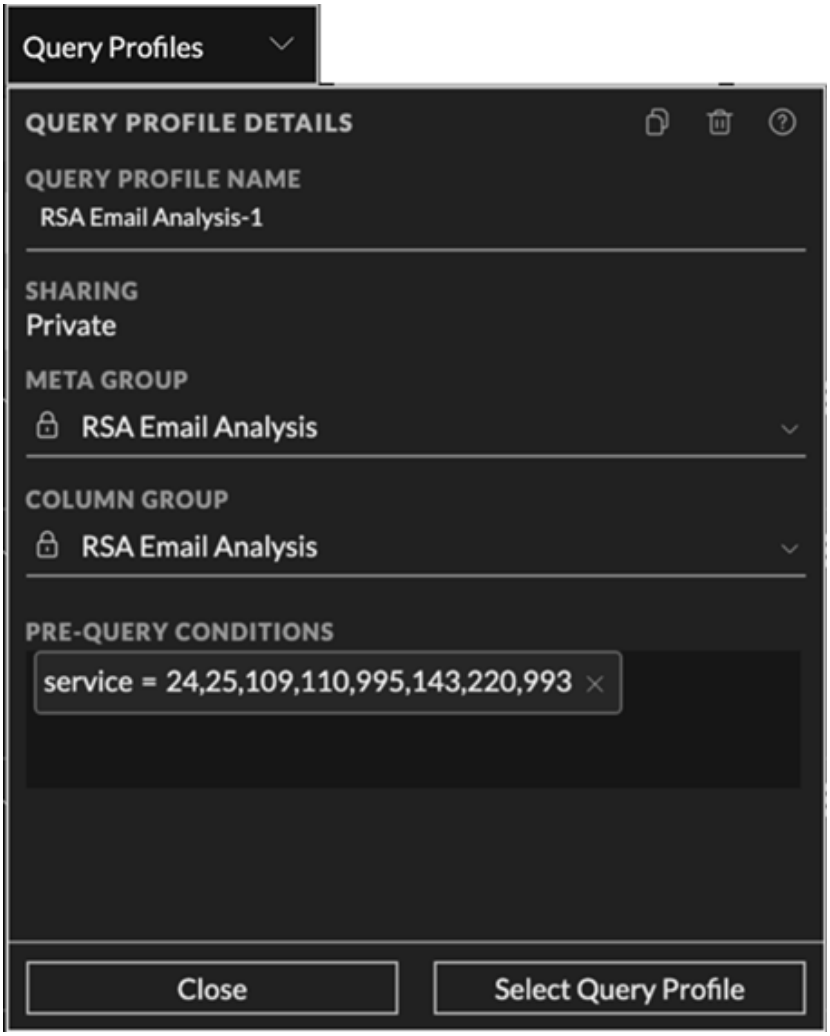
Note: If a Springboard panel is using a query profile as a filter, the profile can be edited, but cannot be deleted in the Events view. However, nothing prevents deletion of the profile in the Navigate view or the Legacy Events view. In this case, Springboard panels that use the deleted query profile as a filter continue to work, but the filter is removed and unexpected results may be displayed in the panel. Refer to "Managing the Springboard" in the *NetWitness Platform Getting Started Guide* for details.

To delete a custom query profile:

1. Go to **Investigate > Events** and click **Query Profiles** in the query bar.
   The Query Profiles menu opens with a list of available profiles.



2. Highlight a custom query profile that you want to delete, and click the edit ( )icon.
   The Query Profile Details dialog is displayed.



3. Click the delete icon ( ).
   In Version 11.5, a confirmation message gives you the opportunity to confirm or cancel the deletion. Click **Cancel** or **Delete Query Profile**.
   In Version 11.4, if the query profile is not a built-in profile, there is no request for confirmation. The profile is deleted and removed from the Query Profiles menu. The profile no longer appears anywhere for any analyst working in Investigate.

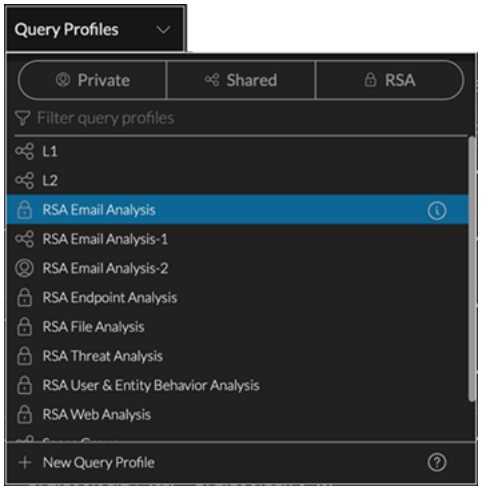
Copy a Query Profile (Version 11.5 and Later)

You can copy any query profile, built-in or custom, shared or private, as long as it does not have unsaved edits in progress. This is useful when you want a customized version of a built-in profile. Also since you cannot change a custom profile from private to shared or from shared to private, creating a copy allows you to select a different Sharing setting. When you copy a profile, the same name is used with a number appended. For example, if you copy RSA Email Analysis, the first copy is named RSA Email Analysis-1, and a second copy of the same profile is named RSA Email Analysis-2. After you create the copy, you can edit the new profile to give it a new name and edit the pre-query conditions, meta group, and column group in the profile.

Note: If you are making a shared copy of a private query profile that uses a private meta group or column group, a message notifies you that a shared copy of the meta group or column group is being created and used in the query profile. It may take a little longer to copy the query profile when a private meta group or column group has to be copied.

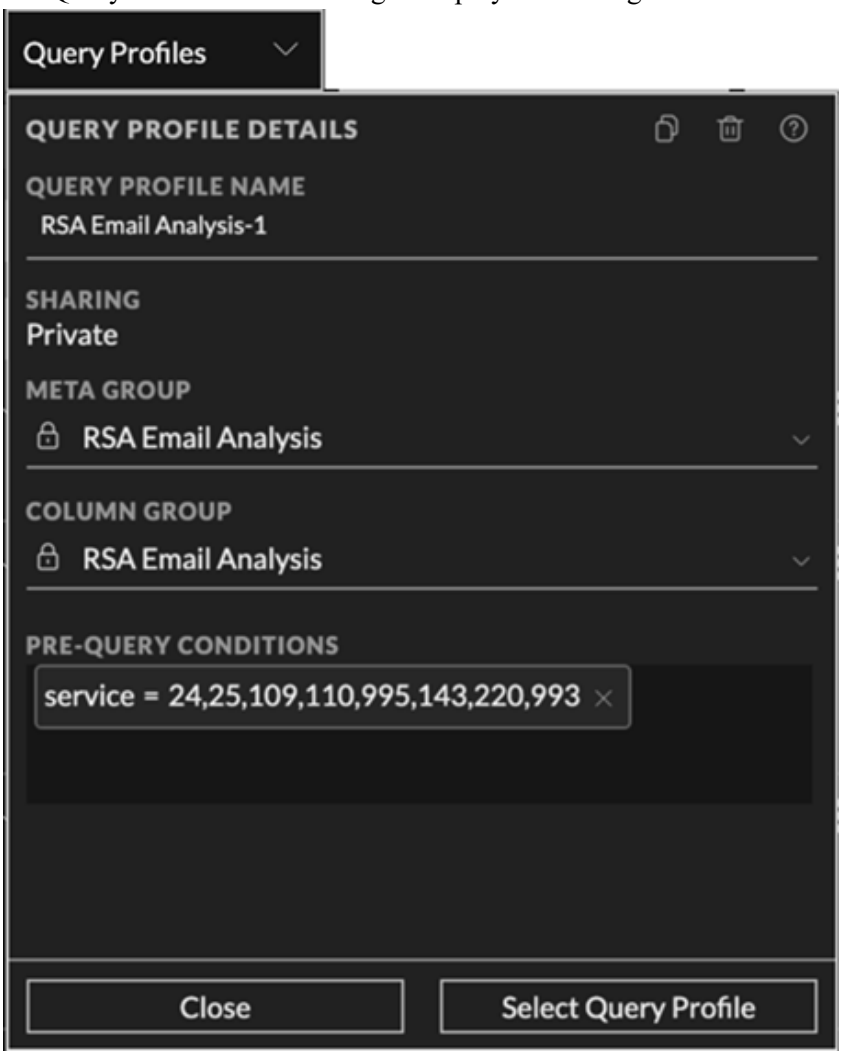
To copy a query profile:

1. Go to **Investigate > Events** and click **Query Profiles** in the query bar.
   The Query Profiles menu opens with a list of available profiles.
2. Highlight the query profile that you want copy. This figure shows RSA Email Analysis highlighted. The information icon ( ) is displayed to the right.

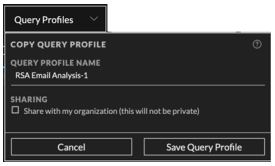


3. Do one of the following:
   a. Click the information icon ( ).
   b. For a custom profile, click the edit icon ( ).

The Query Profile Details dialog is displayed. This figure shows the dialog for a built-in profile.



4. Click the Copy icon ( ).
   The Copy Query Profile dialog is displayed with a number appended to the profile name to create a unique name among all query profiles.



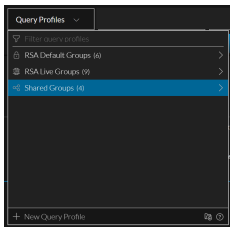
5. (Optional) In the **Query Profile Name** field, edit the name of the query profile.
6. If you want to share the new profile with your organization, set the **Share with my organization** option. By default the new profile is private. If the profile being copied has a private column group or meta group, a shared copy is created and used in the copy of the profile.

7. Do one of the following:
   a. To close the dialog without copying the profile, click **Cancel**.
   b. To save the clone of the query profile, click **Save Query Profile**.
      The clone is saved, and the Query Profile Details dialog for the cloned profile is displayed.
8. Do one of the following:
   a. To close the dialog, click **Close**.
   b. To close the dialog and select the new profile, click **Select Query Profile**.
      The clone is added to the Query Profiles menu.

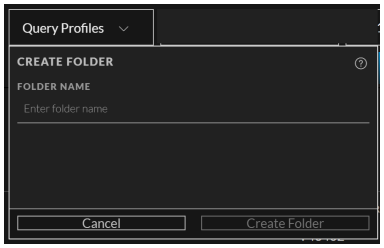
Create a Query Profile Folder

You can create query profiles folders which reside at the top level and are be added as a private or shared folders. And, if the folder name already exists then you are prompted to provide a unique name.

1. In the Events view, select the Query Profiles menu title. The menu drops down to display a list of meta groups and folders with the Filter Query Profiles field at the top and the  option at the bottom.



2. Click  .
   The Create Folder dialog is displayed.

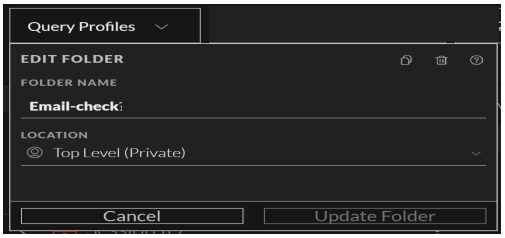


3. In the **Folder Name** field, type a unique name for the new query profile group folder.
4. Click **Create Folder**.

Edit and Move Query Profile Folder

After you create a query profile group folder you can edit or move it, however the folders inside RSA Groups (RSA Live content and RSA OOTB Groups) cannot be edited and moved. The folders inside private and shared folders can be edited and moved only within their respective groups. For example, you cannot move a shared folder into a private folder and vice-versa.

1. In the Events view, select the Query Profiles menu title that you want edit.
2. Click .
   The Edit Folder dialog is displayed.

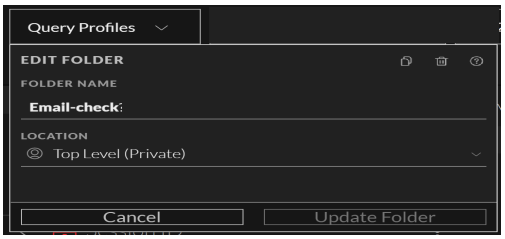


3. In the **Folder Name** field, type a unique name for the query profile folder.
4. Select the location of the folder to be edited.
5. Click **Update Folder**.

Copy Query Profile Folder

You can copy query profiles folder from private to shared, private to private, shared to shared and shared to private groups. When you copy a folder the content inside it gets copied except for the sub-folders. When you copy a private folder into a shared folder, the folder and its content no longer remain private.

1. In the Events view, click the Query Profiles menu title. The menu drops down to display a list of query profiles and folders.
2. Select a folder you want to copy.
3. Click edit and then click .



   The Copy Folder dialog is displayed.

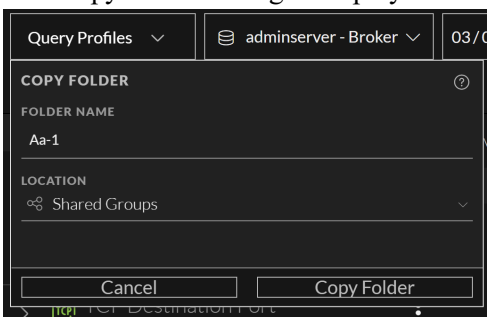


4. In the **Folder Name** field, type a unique name (maximum length of 80 characters) for the new query profile group and folder.
5. Select the location of the folder to be edited.
6. Click **Copy Folder**.

Copy Query Profiles Group Folder Deployed from Live

You can copy query profiles group folder deployed from Live located under RSA Groups category to any other location like Shared groups or to a private folder.

1. In the Events view, click Query Profiles Group menu title. The menu drops down to display a list of query profiles groups and folders.
2. Click on a Live Query Profiles Group folder you want to copy.
3. Click 
   The Copy Folder dialog is displayed.

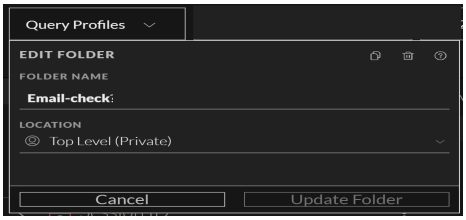


4. Select the location of the folder to be copied.
5. Click **Copy Folder**.
   The folder is created with the original name of the folder and its contents are displayed as the original meta group name appended with a -n.

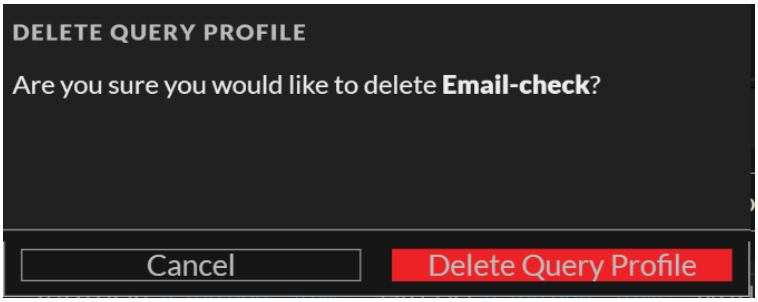
Delete Query Profile Folder

If you don't want to retain a folder you can delete it. However, once the folder is deleted it cannot be retrieved.

1. In the Events view, click the Query Profile menu title. The menu drops down to display a list of query profile groups and folders.
2. Select a folder to be deleted.
3. Click edit  .
   The Edit Folder dialog is displayed.



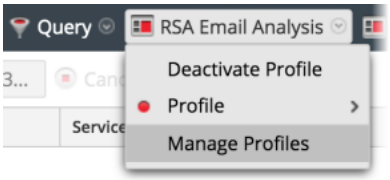
4. Click delete .
   A warning message is displayed to confirm the action.



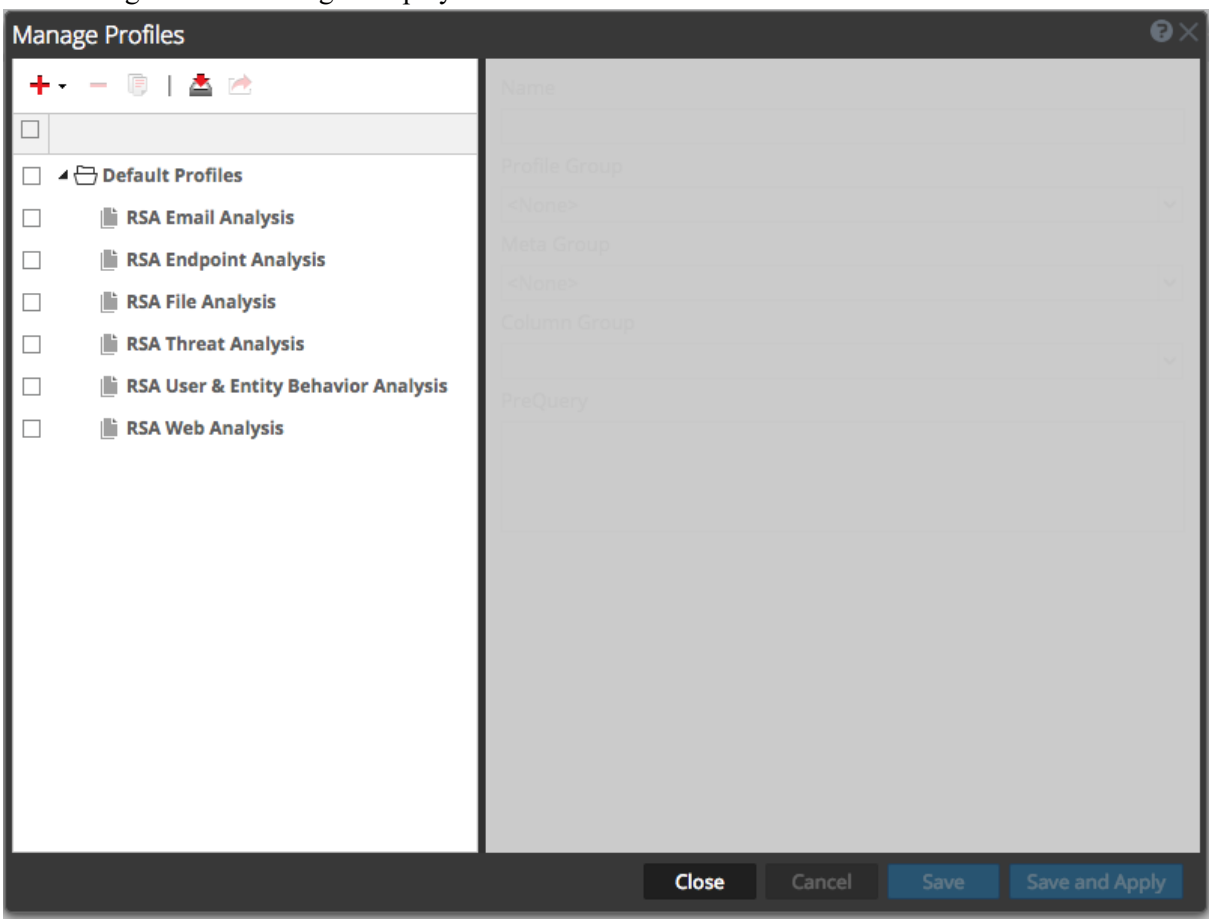
5. (Optional) Select the checkbox, if you want to delete the folder along with all the contents inside the selected folder.
   If you do not select the checkbox, then the content will be moved to the parent folder after the required folder is deleted.
6. Click **OK** to delete.

Navigate to the Manage Profiles Dialog (Navigate and Legacy Events Views)

1. Go to **Investigate > Navigate** or **Legacy Events**. (If the **Investigate** dialog is displayed, select a service and click **Navigate**.)
2. In the toolbar, select **Profile > Manage Profiles**.



The Manage Profiles dialog is displayed.



Create, Edit, or Delete a Profile Group (Navigate or Legacy Events View)

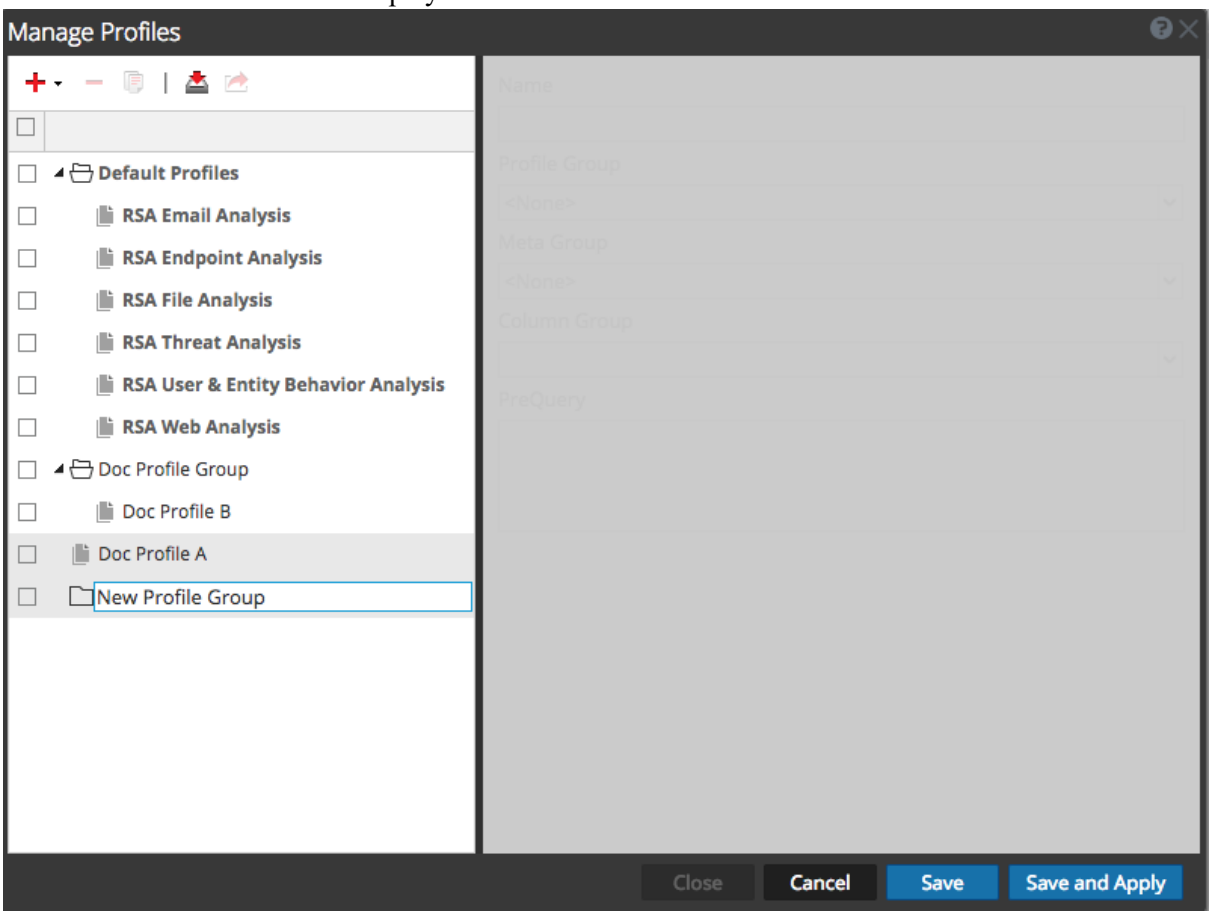
You can create a custom profile group to organize different profiles. Once created, the only edit you can make directly to a profile group is to edit the name of the profile group. To add or remove a profile in a group, edit the profile and assign it to a different profile group as described in Create and Edit Profiles (Navigate or Legacy Events View).

Note: If you migrated profile groups from Version 11.3, empty groups were not migrated.

1. In the **Manage Profiles** dialog, do one of the following:
   - To select an existing profile group to edit, double-click the profile group.
   - To add a new profile group, click  and select **Add New Profile Group**.

Note: If you want to edit one of the built-in profile groups, click  to make an editable copy.

A folder with a blank field is displayed at the bottom of the Profiles list in the left column.



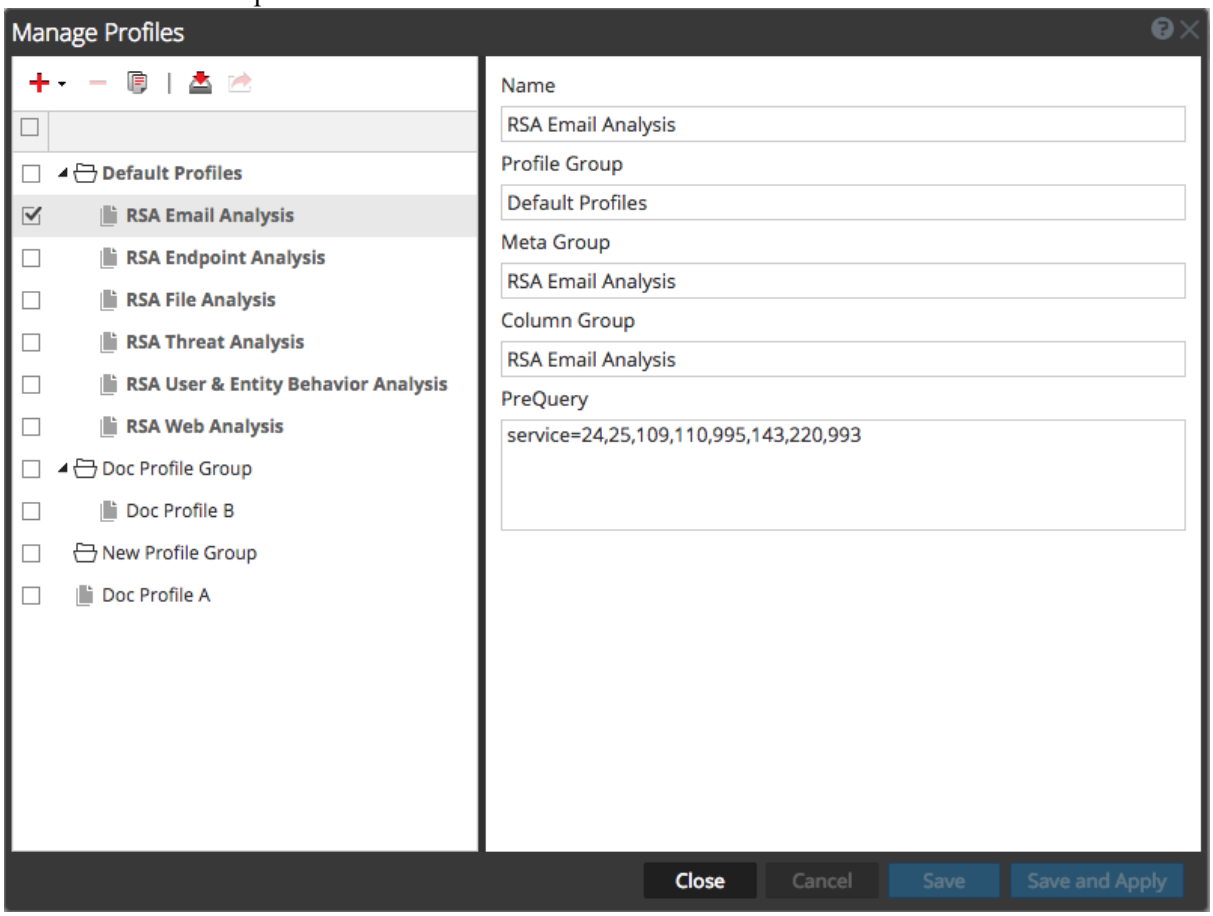
2. To edit or enter the name of the profile group, double-click the Profile Group and type in the entry field. The name must be between 2 and 80 characters.
   The profile group name is applied to a new profile group or to the profile group you edited. The profile group is now available when configuring a profile.
3. To delete a profile group do one of the following:
   - If you want to delete a profile group but keep the profiles, click the checkbox to select the group, uncheck the profiles in the group, and click delete.
   - If you want to delete a profile group and the profiles that the group contains, click the checkbox to select the group, and leave the profiles that you want to delete checked.
     A dialog asks for confirmation that you want to delete the group. If you left the mark in the checkbox next to the profiles, the group and the profiles in the group are deleted. If you unchecked the profiles, only the profile group is deleted and the profiles are moved out of the group and available to add to another profile group.

Create and Edit Profiles (Navigate or Legacy Events View)

1. In the **Manage Profiles** dialog, do one of the following:
   - To select an existing profile to edit, click the checkbox beside the name.
   - To add a new profile in Version 11.2 and later, click + or click the down arrow next to + and select **Add New Profile**.
   - To create a new profile in versions prior to 11.2, click +.

   Note: If you want to edit one of the built-in profiles, click to create a copy, and edit the copy.

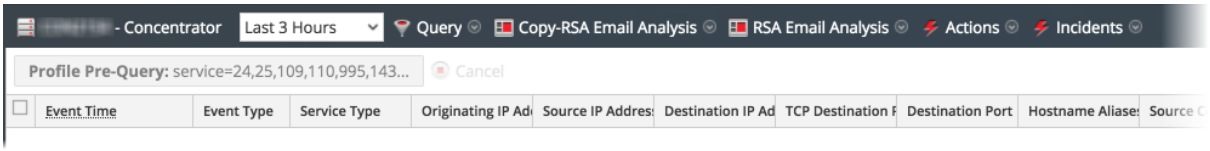
   The definition of the profile is available to edit in the right panel. This figure illustrates the definition of one of the built-in profiles.



2. Edit or enter the profile name by typing in the **Name** field. The name must be between 2 and 80 characters.
3. (Optional for Version 11.2 and later) If you want to add the profile to a profile group, select a profile group from the **Profile Group** drop-down list.

If you select a profile group, the profile is added to the group when you save the changes. If you do not select a profile group, the profile is not part of a group.

4. Select a meta group from the **Meta Group** drop-down list. You can add custom meta groups as described in Use Meta Groups to Focus on Relevant Meta Keys. Private meta groups created in the Events view are not available in the Navigate view.
5. Select a column group for the **Column Group** drop-down list. You can add custom column groups as described in Use Columns and Column Groups in the Events List. Private column groups created in the Events view are not available in the Navigate view.
6. Type queries to filter results in the **PreQuery** field. PreQuery follows the same syntax as the Query builder. The PreQuery in the figure uses a meta group called **service = 24,25,109,110,995,143,220,993**.
7. Click **Save** to save the profile without using it, or click **Save and Apply** to save the profile and use it immediately.
   If you click **Save and Apply**, a confirmation dialog is displayed before applying the selected profile. For Version 11.2 and later, the PreQuery that you entered in the Manage Profiles dialog is displayed in the breadcrumb.



Delete a Profile (Navigate or Legacy Events View)

1. In the **Manage Profiles** dialog, select a profile by clicking the checkbox beside the name.

   Note: You cannot delete any of the built-in profiles.

2. Click −.
   A prompt requests confirmation that you want to delete the profile, and the profile is deleted. The option name in the toolbar reverts to **Profile** to show that no profile is in effect.

Change the Active Profile (Navigate or Legacy Events View)

If you do not see enough results or the right results in the Navigate or Events views, you may have an active profile that is applying a PreQuery. If you do not want to use any profiles, you can click **Deactivate Profile** in the **Profile** drop-down menu.

To use a different profile:

1. In the **Navigate** or **Legacy Events** view toolbar, open the **Profiles** drop-down menu.
2. Hover over the **Profile** option to display a drop-down list of available profiles.
3. Select the profile you want to use.
   The profile settings are applied immediately.

If you want to change the active profile from the Manage Profile dialog:

1. In the **Navigate** or **Legacy Events** view toolbar, select **Profiles > Manage Profiles**.
   The Manage Profiles dialog is displayed.
2. Select a profile from the left panel and click **Save and Apply**.
   A confirmation dialog is displayed.
3. Click **Yes**.
   The profile settings are applied immediately.

Import Profiles (Navigate or Legacy Events View)

In the Navigate view and the Legacy Events view, you can upload or import .jsn files that have been downloaded from another service. When profile groups are exported and then imported, the grouping of profiles is maintained.

1. In the **Manage Profiles** dialog, click  in the left panel toolbar.
   The Profile Import dialog is displayed.
2. Click **Browse** or the **Upload File** field to select a file from your computer.
3. When the file is selected, click **Upload**.
   The profile is displayed in the left panel.

Download Profiles (Navigate or Legacy Events View)

In the Navigate view and the Legacy Events view, profiles are downloaded as .jsn files.

1. In the **Manage Profiles** dialog, select one or more profiles from the left panel.
2. In the left panel toolbar, click  .
   The download begins immediately.

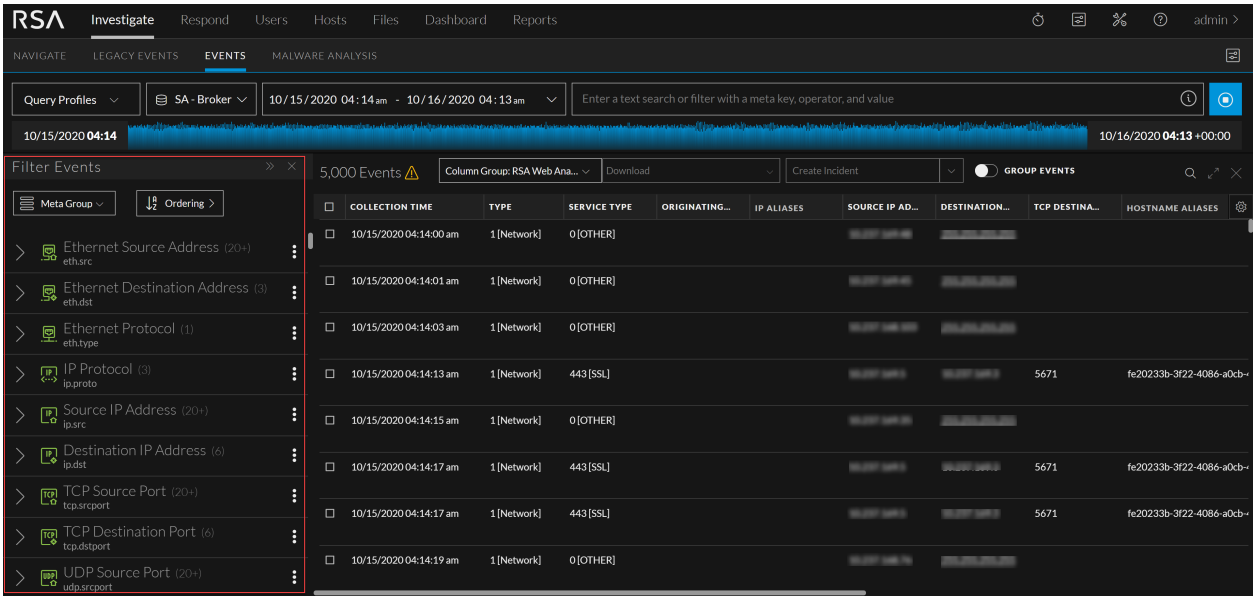
Drill into Metadata in the Events View (BETA)

Note: This section applies to Version 11.5 and later. The feature is a beta feature that is enabled by default, and can be disabled by the system administrator as described in the *System Security and User Management Guide*.

When working in the Events view, the focus of an investigation is the smallest possible set of relevant events in sequential order. You can reduce the number of visible events loaded in the Events view using query profiles, column groups, meta groups, and queries. However, it is more efficient to limit the data set using the metadata indexed on the Concentrator before looking at the actual events stored on the Decoder or Log Decoder.

In Version 11.4.x and earlier, it is best to start by looking at the meta keys and meta values indexed on the Concentrator and drill into the metadata in the Navigate view to find a relevant set of events, with each drill or query further limiting the data set. When you have a meaningful data set, or drill point, you can examine the details of the related events in sequential order in the Events view.

Beginning with Version 11.5, you can drill into the metadata in the Filter Events panel, without leaving the Events view. The list of meta keys and meta values shown is related to all events seen in the environment for the time range in the query. When you find the drill point of interest in the Filter Events panel, you can open the Events panel to see the sequential events. The set of events loaded in the Events view is smaller and loads faster. The flow of an investigation is smoother with less hopping between views. The figure below illustrates the panel, open to the left of the Events panel.



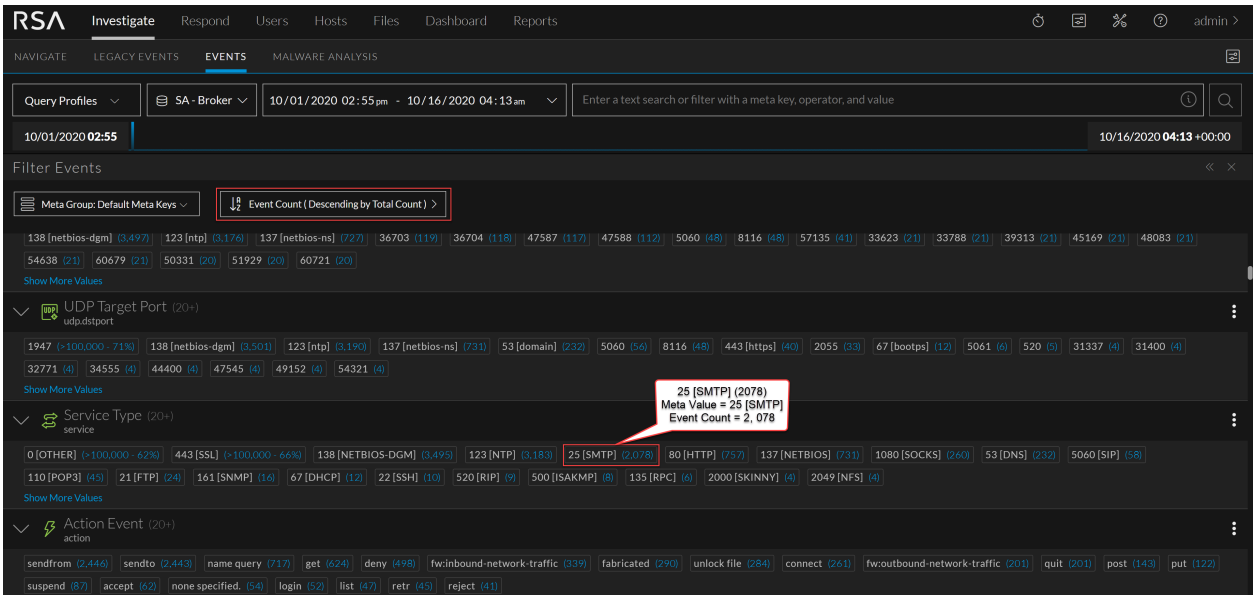
Note: There are two situations in which results in the Filter Events panel may not be as expected:
-In a mixed-mode environment with a Version 11.5 Broker and some Core services at RSA NetWitness Platform Version 11.4 or earlier, a text filter is not supported in the Filter Events panel. If the query in the Events panel includes a text filter, the result set in the Events panel and Filter Events panel may be different.
-If the query in the Events view query builder has a logical `OR` or `&&`, the results in the Events view may be different from results for the same query in the Navigate view and Legacy Events view. In this situation, a set of parentheses automatically encloses the logical `OR` expression in the Navigate view and Legacy Events view, while parentheses have to be manually added in the Events view. If this occurs, you need to enclose the logical OR expression in an additional set of parentheses; select the two filters in the query bar, right-click one of them, and select **Wrap in parentheses** in the menu.

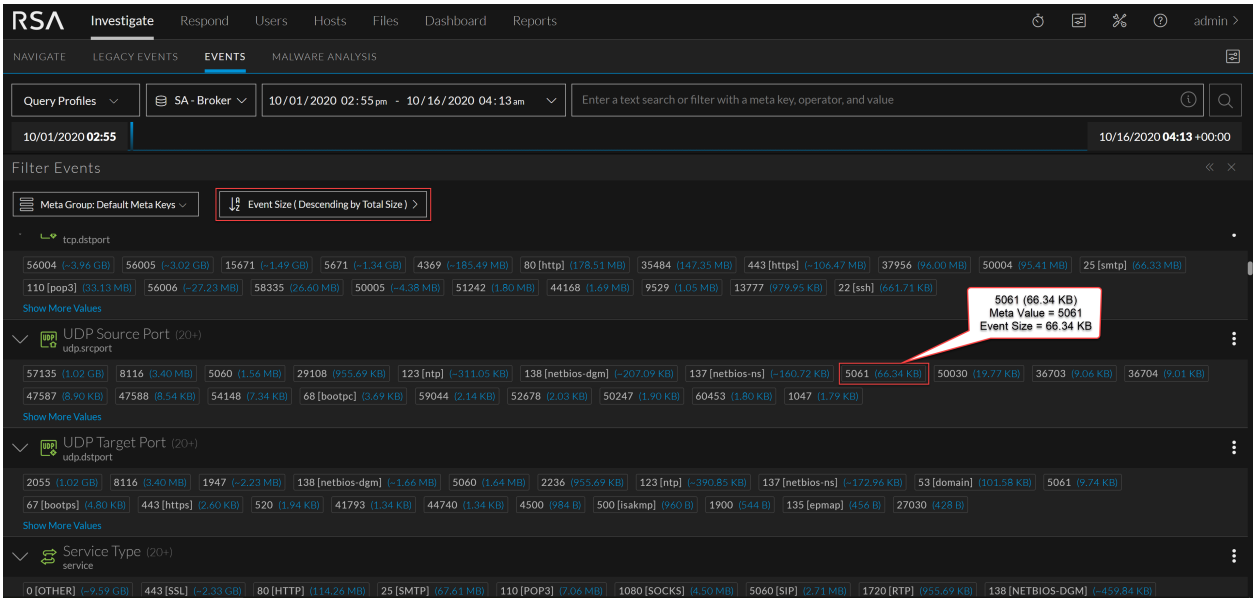
Note: (Version 11.5.1) In the Filter Events panel, the meta values result threshold is 100000. If results are above the threshold, it is indicated using either ~ or >. For example, **(>100000)** indicates that the results are sorted based on count and are greater than the threshold. Similarly, **(~100000)** indicates that the results are sorted based on size and are greater than the threshold.

Modes of Operation

The Filter Events panel has two modes of operation.

- The **narrowFilter Events panel** is part of a faceted search view into the data (shown above). Left- or right-clicking a meta value adds a new filter, automatically executes a new query, and displays matching events in the sequential list of events. When both panels are open, you can drill into the data in both the Filter Events panel and the Events panel . Each time you left-click a meta value in the Filter Events panel, an expression is appended to the query bar, and the query is executed by default. The query results show new metadata to filter by in the Filter Events panel and the resulting events that match the query in the Events panel. If you change the service or other query elements in the Events panel, you need to execute the query to reload the Filter Events panel.
- The **fully expanded Filter Events panel** uses the full width of the browser window to provide ample real estate to hunt through the metadata without the performance load of immediately submitting a query or viewing the sequential events. As you click a new meta value and drill into the meta values, each meta value is added to the query filter and executed in the Filter Events panel, so that the number of events seen is reduced. Because the Events panel is closed, the query in the Events panel is not updated and the query is not executed. When you collapse the Filter Events panel back to original size, the Events list opens and the query is executed. This is an example of the fully expanded panel.



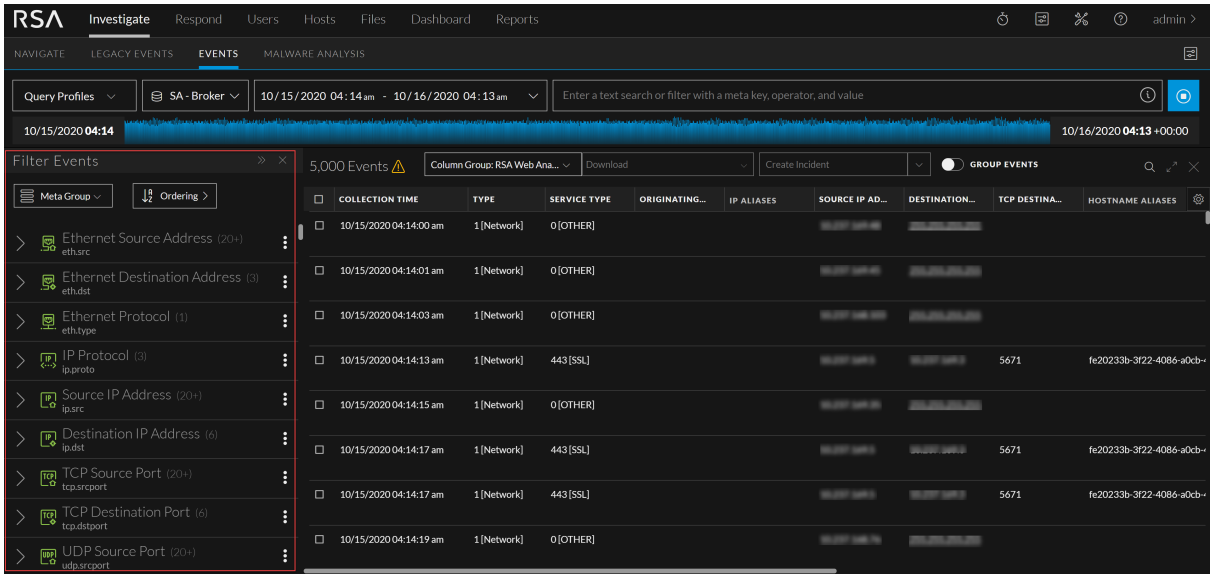


View Metadata in the Filter Events Panel

To view metadata in the Filter Events panel:

1. Go to **Investigate > Events**, select a service to investigate, and select a time range.
2. (Optional) Select a column group or a query profile.
3. Click  to load events in the Events panel.
   A query is executed in the Events panel and matching events are listed,

4. Click the Filter button (Filter) in the Events panel.
   The Filter Events panel opens to the left of the Events panel.

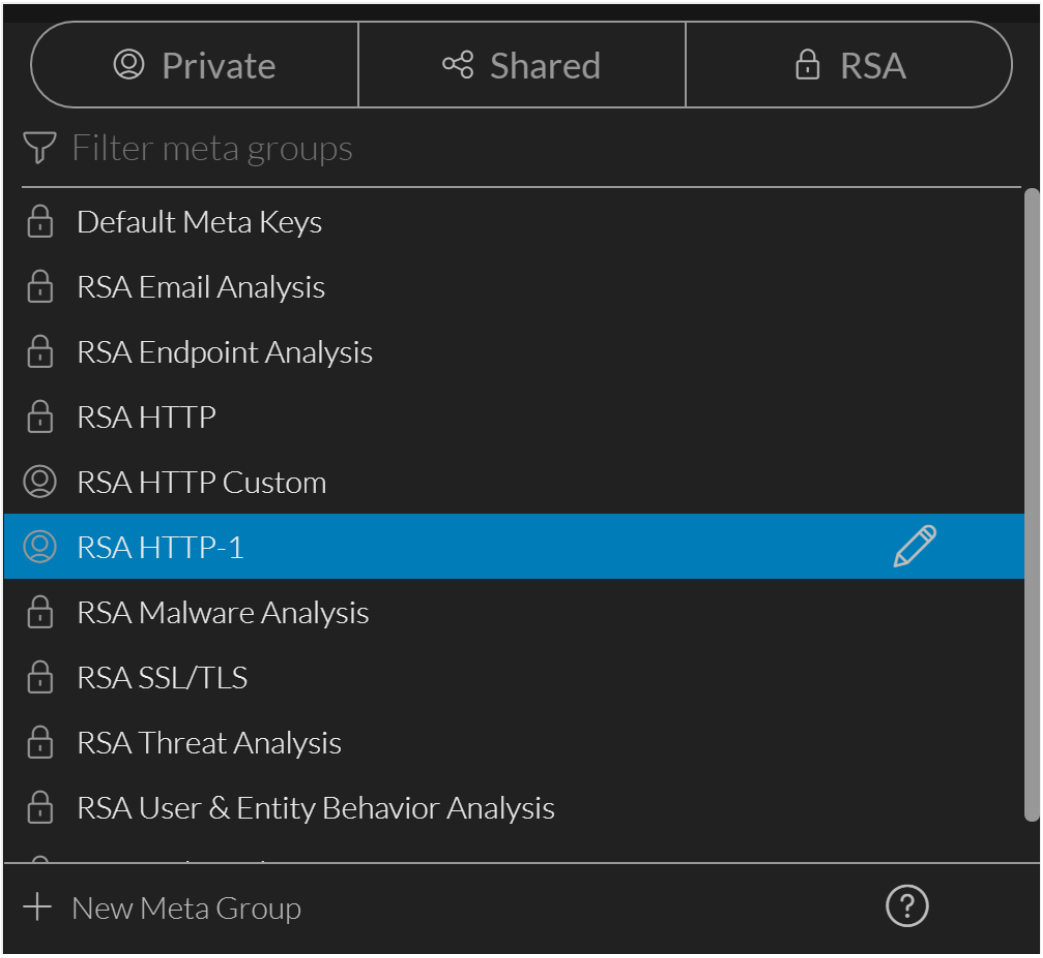


Note: (Version 11.6) By default, the Filter Events panel is open in the Events view. The last used state of the panel (narrow or fully expanded) is saved throughout the session and across logins. Also, the Filter Events panel provides additional contrast between meta keys, meta values, and meta counts to improve readability.

The Default Meta Keys meta group is in effect the first time you log in. If you selected a different meta group the last time you logged in, it remains in effect until browser cache is cleared. In Version 11.5.1, the meta group you selected previously is not stored in browser cache so it remains in effect until you change it. See Use Meta Groups to Focus on Relevant Meta Keys for details about meta groups. Based on the contents of the index file for the service, the Filter Events panel is populated with the first 25 meta keys that have at least one meta value and are open. When using the Default Meta Keys group in the Filter Events panel, only the first 30 meta keys with values are open and the remaining are closed. Closed meta keys may be listed, but they do not count toward the 25 or 30 meta keys total. Meta keys with no values are listed at the bottom of the panel. You can expand, collapse, and close the panel using the standard panel controls (, , and).

1. Do one of the following:
   a. To close the dialog without editing, click **Close**.
   b. To close the dialog and select the copy of the meta group, click **Select Meta Group**.
      The group is added to the Meta Group menu. The figure below has a private copy of the RSA

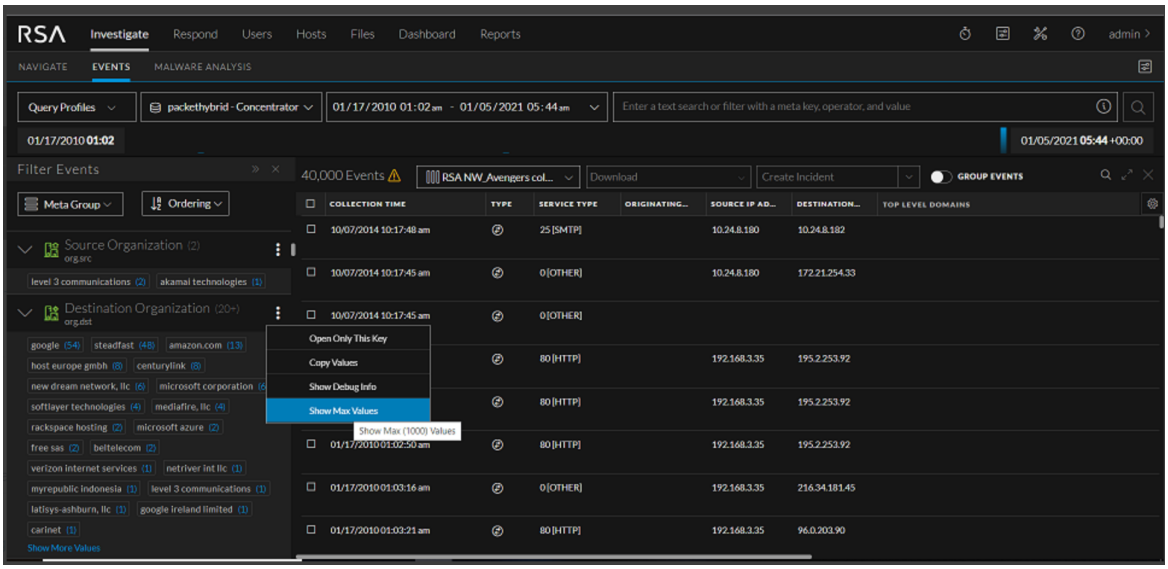
HTTP meta group.



Show Max Value of Meta Groups

In case all the values have not rendered and displayed, you can click Show Max Value to view all the values at once.

1. With the Filter Events panel open in the 11.6 Events view, click  and select the **Show Max Values** option.



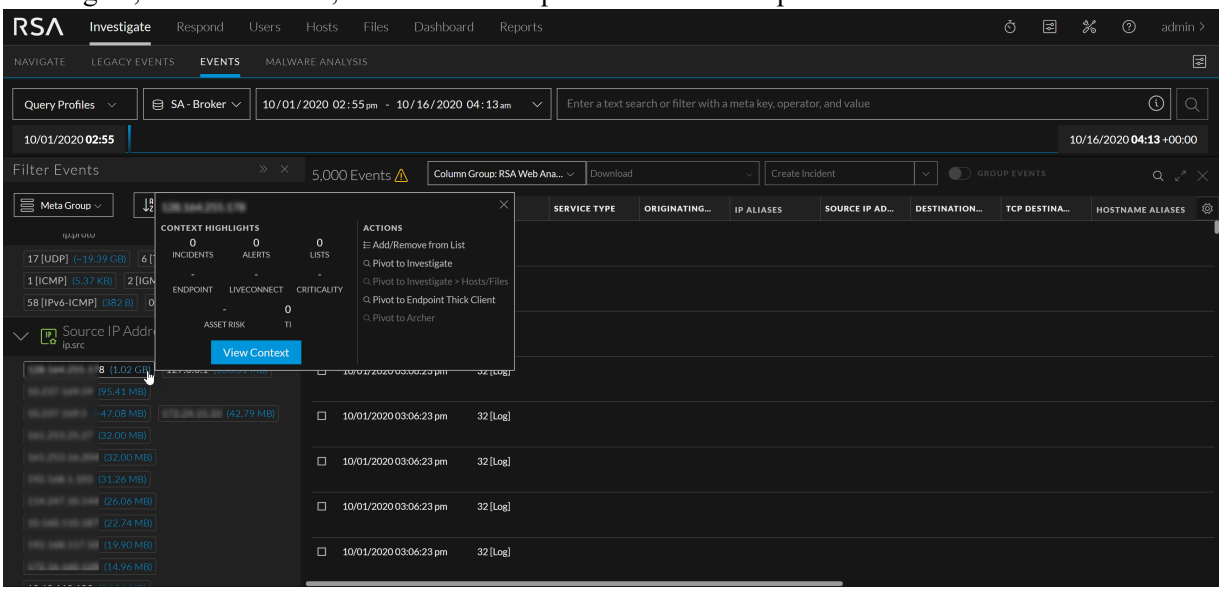
2. The values that were not rendered earlier will begin to load and a maximum of 1000 results are displayed.

View the Context Lookup Panel in the Filter Events Panel

In the Filter Events panel, you can click a meta entity to open the context tooltip. The context tooltip is available only for the meta keys that are defined as an entity the Context Hub supports. The Context Hub service is pre-configured with default meta types and meta keys mapping. For information about mapping of the context hub meta values with investigation meta keys, see "Configure Meta Type Mapping for Context Hub" in the *Context Hub Configuration Guide*. The context tooltip includes the following two sections.

- **Context Highlights** - The information in this section helps you to determine the actions that you would like to take. It can show related data for Incidents, Alerts, Lists, Endpoint, Live Connect, Criticality, Asset Risk, and Threat Intelligence (TI). Depending on your data, you may be able to click these items for more information.
- **Actions** - This section lists the available actions. In the example, the Add/Remove from List, Pivot to

Investigate, Pivot to Archer, and Pivot to Endpoint Thick Client options are available.



Understand Visible Metadata

Each meta key has a list of meta values, with up to 20 values displayed by default. You can click **Show More Values** to incrementally add 20 meta values, up to a total of 1,000 meta values, which is a hard-coded limit to optimize performance. The meta key name and plain English name of each meta key found in the service, both populated and non-populated, are listed. For each meta value, you can see the number of events in the current results that contain the value (count) or the size of the events in the current results (size). For example, the following might be listed:

```
Action Event [action] (3)
get(3016) login (1346) put (501)
```

In this example, the meta key name is `action`, the English name is Action Event, and three meta values were found for this meta key. There were 3016 events containing `get`, 1346 events containing `login`, and 501 events containing `put`. The values are ordered so that the value with the largest count is listed first.

In the following example, the same meta key has the values ordered based on the event size in bytes. The smallest size is listed first:

```
Action Event [action] (3)
login (13,034,588) put (21,848,760) get (1,409,079,256)
```

An icon before each meta key name identifies the indexing method for the key. The indexing method determines the types of interactions and queries possible using that meta key.

- This meta key is indexed by value: Action Event [action] (40+). The green color indicates that the all available interactions and queries are supported. You can see the available interactions in the context menu by right-clicking the meta value.
- This meta key is indexed by meta key: Bytes Sent [bytes.src]. The yellow color is a clue that a subset of available interactions is supported, and queries on this meta key may take longer than meta keys that

are indexed by value. You can see the available interactions in the context menu by right-clicking the meta value.

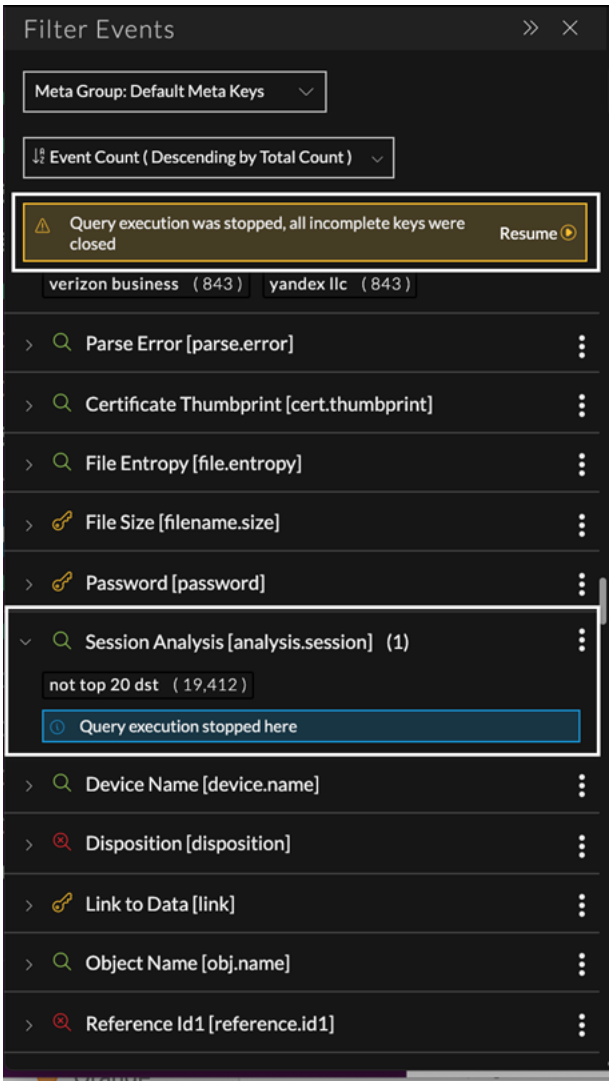
- This meta key is not indexed: MAC Alias Record [alias.mac]. Values for non-indexed meta keys cannot be used to query. If you want to query a meta key that is not indexed, your administrator needs to edit the index file for the service to index the meta key by value or meta key.
- For Version 11.5.1, a set of more than 200 meta key symbols replaces the three indexing method symbols to provide a visual indicator of the purpose of the meta key. The color of the meta key symbol identifies the indexing method using the same colors as before: green, yellow, and red. A tooltip also identifies the indexing method and provides a description of the icon. The icons are defined based on categories outlined in the Unified Data Model (https://community.rsa.com/community/products/netwitness/rsa-content/udm). There is a generic icon for most categories that do not have specific meta keys and a default meta key icon to use when a new custom meta key is added.

If an error occurs while loading a meta key, the other meta keys load as usual and an error message is displayed in the meta key that did not load. When you execute a new query, some error messages disappear. Meta keys that have no values in the set of events are listed at the bottom of the panel.

Stop and Resume Metadata Loading

You can stop and resume loading of metadata in the Version 11.5.1 Filter Events panel as the meta keys and values are loading. When loading a lot of metadata this can save time because you do not have to wait for all data to load. If you stop loading and still need to see more metadata, you can resume loading, then stop again when you see the data you want.

1. While data is loading in the Filter Events panel, click the Stop button ( ) in the query bar.
   Meta keys stop loading, and keys that did not finish loading are closed. A message above the meta keys list informs about the status, and you can scroll down to find the last meta key to finish loading. In this example, Session Analysis finished loading; all the meta keys below remain closed.



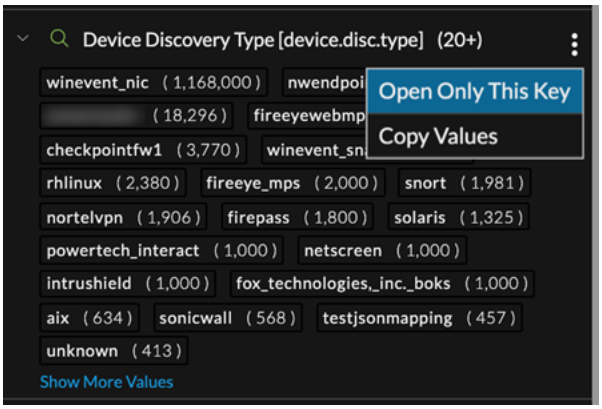
2. Do one of the following:
   a. To resume loading at the meta key where loading stopped, click **Resume**.
   b. If you want to review the values of specific meta keys that were not loaded, without resuming the query, click the meta key name to open any key.

Close All Except One Meta Key (Version 11.5.1 and Later)

When the Filter Events panel is open, seeing many meta keys at the same time can be distracting.

To close all except one meta key:

1. In the meta key row of an entry, click the Meta Key options button (⋮).
   The Meta Key options are displayed.



2. Select Open Only This Key.
   All except the current meta key close. If the selected key is closed, it opens and data is loaded, while all other keys close.

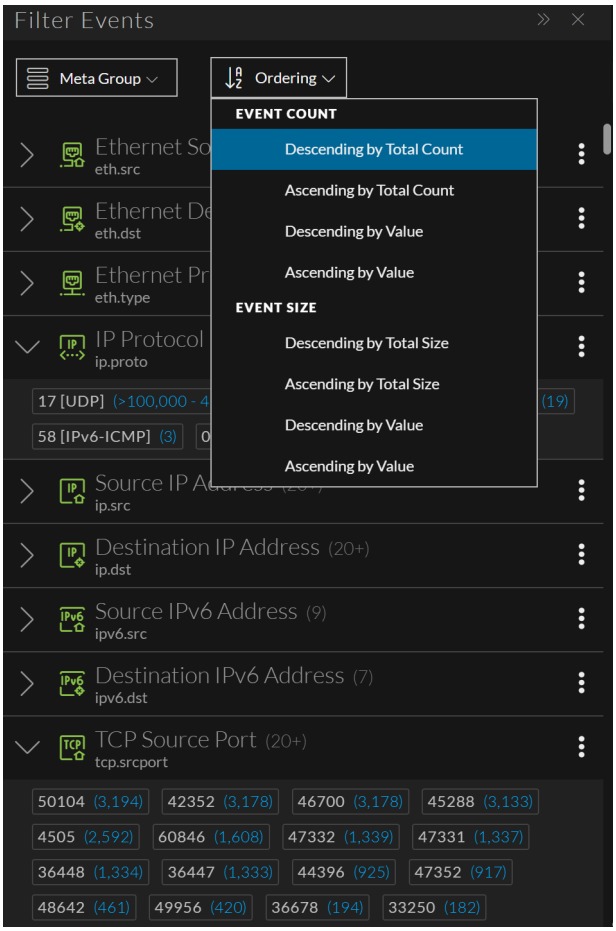
Set the Ordering Method for Meta Values

With the Filter Events panel open, you can look at two parameters for each value: the event count or the event size. Each meta key entry includes either the event count or the event size in parentheses after the value. In both cases, there are four options for ordering.

To use the ordering options:

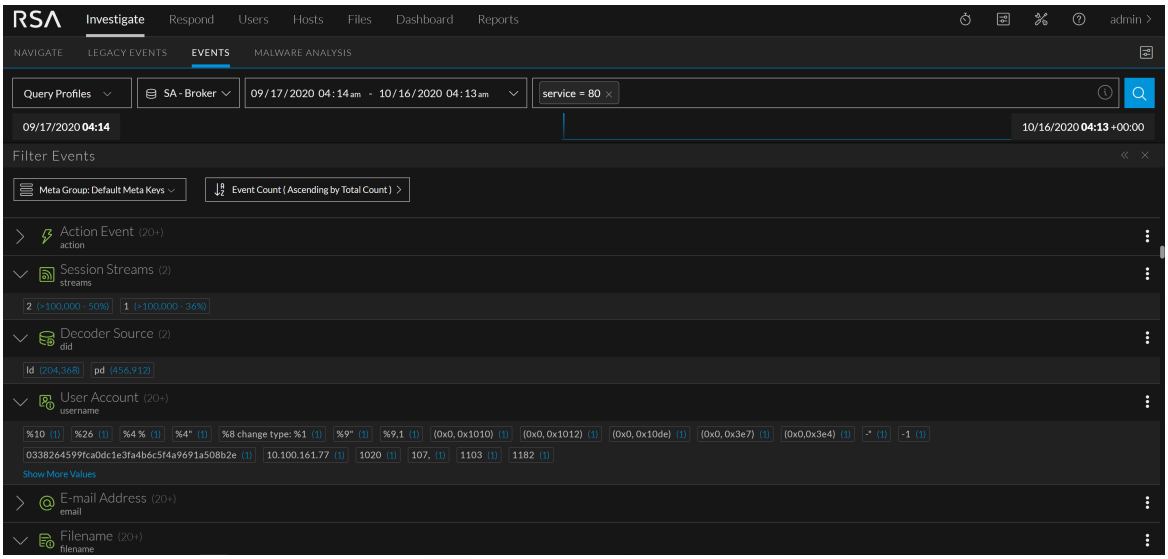
1. With the Filter Events panel open, click the ordering menu label, which is named according to the selected ordering option. This is an example of the menu label when ordering by event count in ascending order by total count: 

   .
   The Ordering menu is displayed. This figure shows the narrow version of the menu.



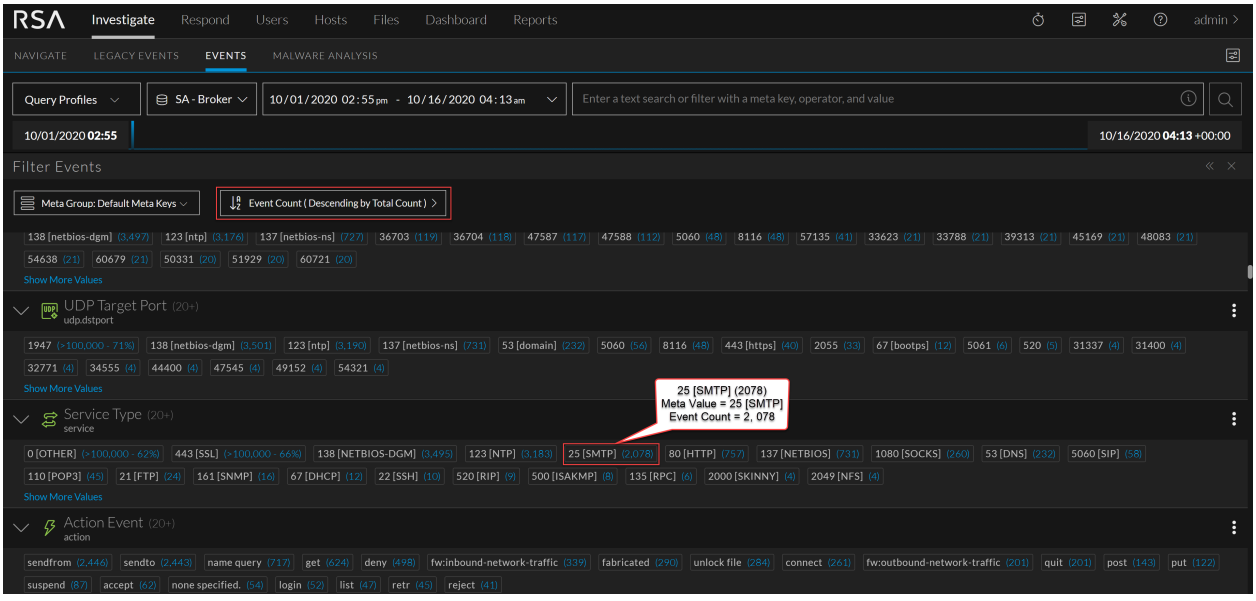
2. If you want to see the event count in parentheses after each value, select one of the following options. By default, the meta keys are displayed using the Event Count > Descending by Total Count method.
   a. To order by total count of events in which the value was found, select either **Descending by Total Count** or **Ascending by Total Count**.
   b. To order by the name of the value, select either **Ascending by Value** or **Descending by Value**.
3. If you want to see the size in bytes of the events in which the value was found, select one of the following options.
   a. To order by total size of events in which the value was found, select either **Descending by Total Size** or **Ascending by Total Size**.
   b. To order by the name of the value, select either **Ascending by Total Size** or **Descending by Total Size**.
   Under each meta key in the Filter Events panel, the values are ordered according to your

selection.

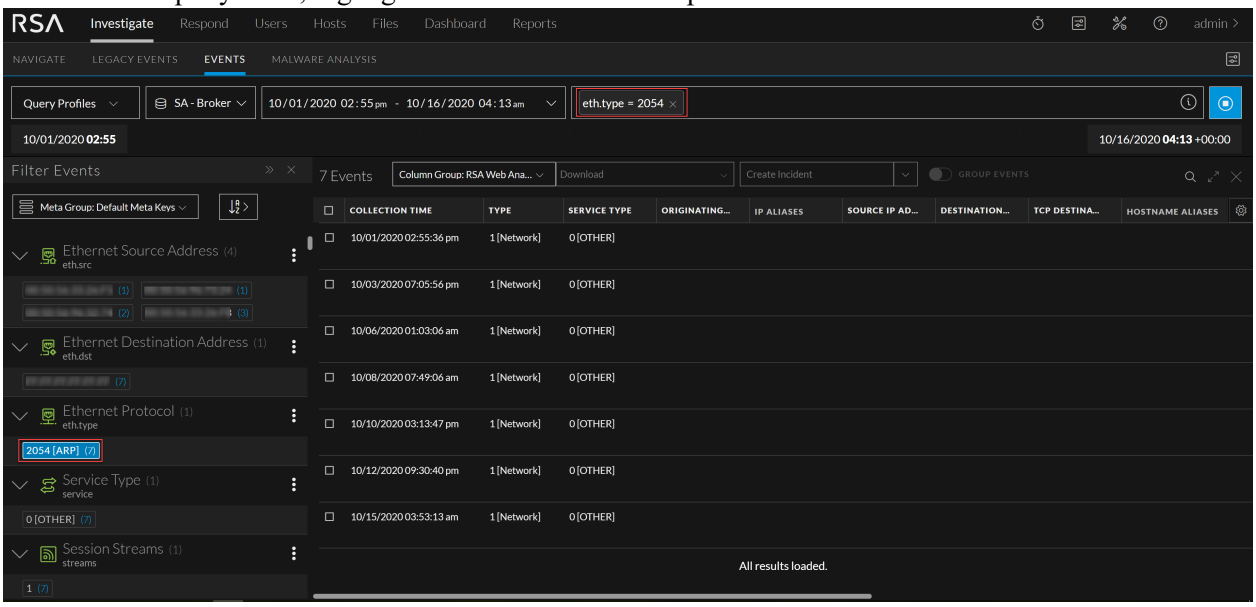


Drill into Meta Values

With the Filter Events panel open, you can drill into meta values to focus an investigation down to the smallest possible set of relevant events. Drilling in the fully expanded Filter Events panel adds filters to the query bar and refines the displayed metadata in the Filter Events panel, but does not execute the query in the Events panel. Drilling in the narrow panel, side by side with the Events panel, adds the filter to the query bar and executes the query in the Events panel and the Filter Events panel. This figure is an example of the fully expanded panel with some metadata loaded.



You can drill into metadata in the Filter Events panel to find relevant meta values. A simple query using the (=) operator highlights the meta value used in the Filter Events panel. This helps to associate the metadata with the filter added to the query. For example, the following figure shows the meta key value, related to the query filter, highlighted in the Filter Events panel.



To drill into meta values in the fully expanded Filter Events panel:

1. Look for a meta value that is of interest, and click the value. Using the figure above as an example, to investigate the SMTP service type as opposed to other service types, click **25[SMTP]**.
   The other service types are filtered out of the metadata in the Filter Events panel, but the query is not executed in the Events panel.
2. Look for a meta value that is of interest, and do one of the following:
   a. Click the value. Using the figure above as an example, to investigate the SMTP service type as opposed to other service types, click **25[SMTP]**.
      The filter is added as the last filter in the query bar, and other service types are filtered out of the metadata in the Filter Events panel. With the Events panel closed, no query is executed there.
   b. (Version 11.5.1) Right-click the value and select **Add Filter - Do Not Run Query** in the drop-down menu.
      The filter is added as the last filter in the query bar, but no other service types are filtered out of the metadata in the Filter Events panel. With the Events panel closed, no query is executed there.
   c. (Version 11.5.1) Press **CTRL (Windows)** or **CMD (MacOS)** and click the value.
      The filter is added as the last filter in the query bar, but no other service types are filtered out of the metadata in the Filter Events panel. With the Events panel closed, no query is executed there.
3. Repeat step 1 with another meta value, for example, **writetoexecutable** in the **Action Event [action]** meta key. Continue drilling into values until you find a set of events (drill point) that you want to see in sequential order.

4. To view the sequential events for the drill point, click  to shrink the Filter Events panel.
   The Events panel opens to the right, and the query is executed in the Events panel so that you can see the raw events in sequential order.

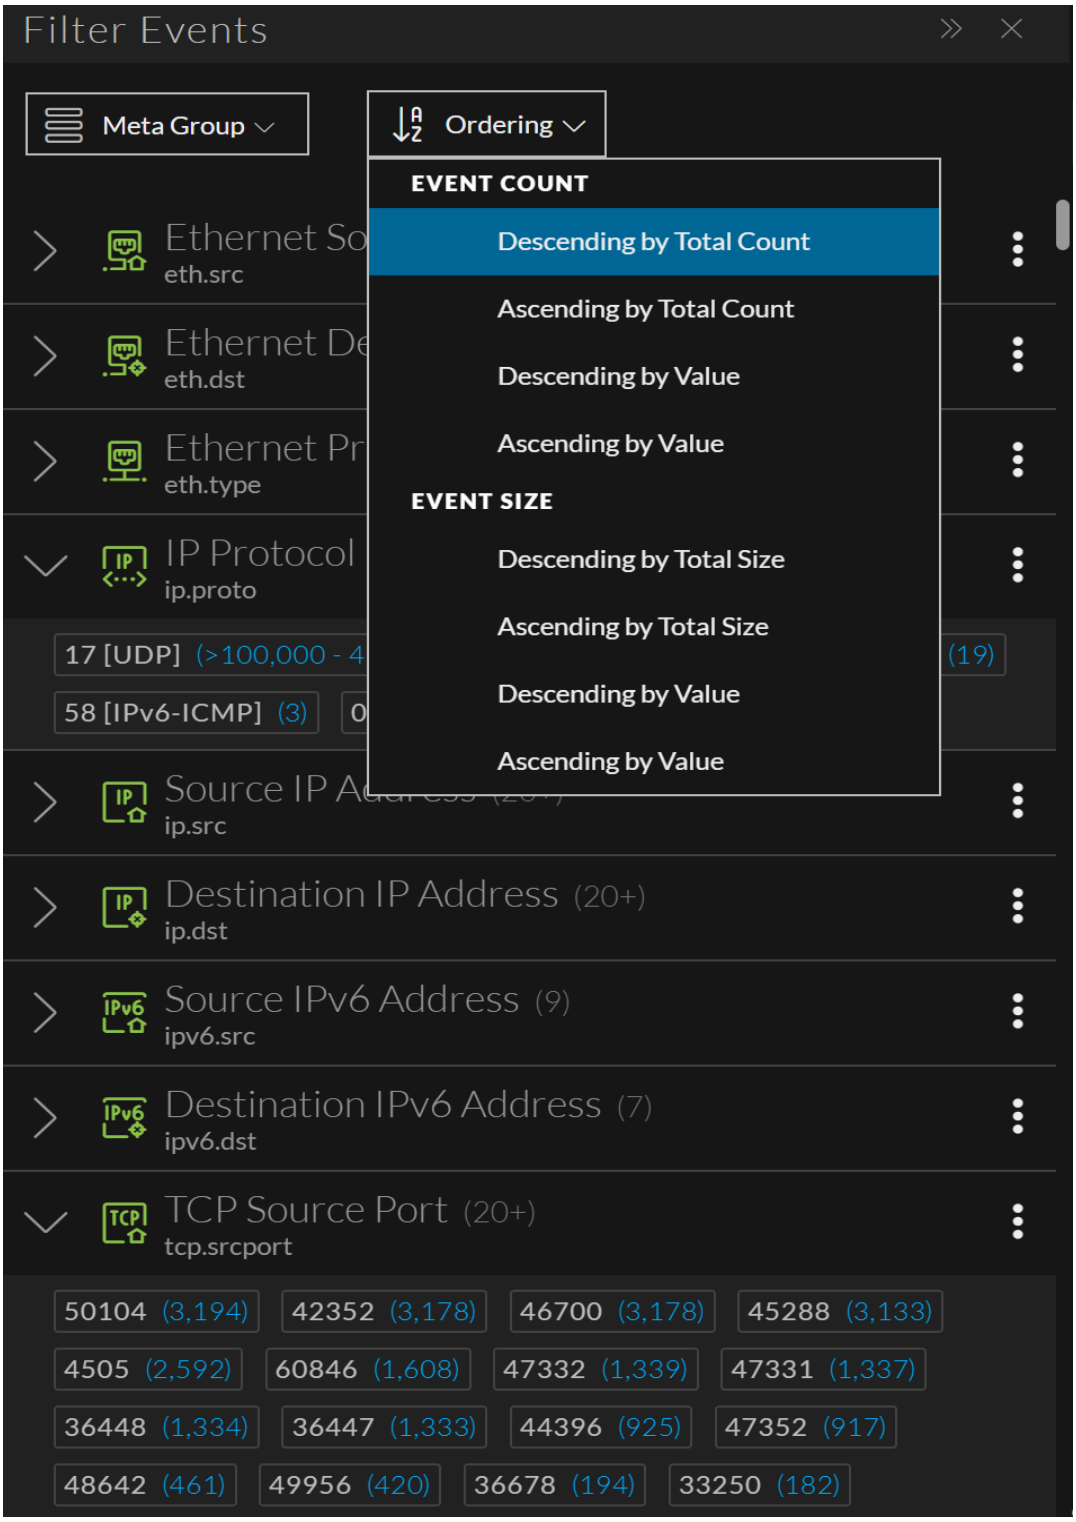
To drill into meta values in the narrow Filter Events panel:

1. Look for a meta value that is of interest, and do one of the following:
   a. Click the value. Using the figure above as an example, to investigate the SMTP service type as opposed to other service types, click **25[SMTP]**.
      The filter is added as the last filter in the query bar, and other service types are filtered out of the metadata in the Filter Events panel and the data set showing in the Events panel.
   b. Right-click the value and select **Add Filter - Do Not Run Query** in the drop-down menu.
      The filter is added as the last filter in the query bar, but no other service types are filtered out of the metadata in the Filter Events panel, and the query is not executed in the Events panel until you click the query button.
   c. Press **CTRL (Windows)** or **CMD (MacOS)** and click the value.
      The filter is added as the last filter in the query bar, but no other service types are filtered out of the metadata, and the query is not executed in the Events panel until you click the query button.
2. Continue clicking values to refine the set of events (drill point). As you refine the set of events, examine and reconstruct the raw events for the same set in the Events panel.

Copy the Meta Values for a Meta Key

To copy all of the visible meta values for a meta key:

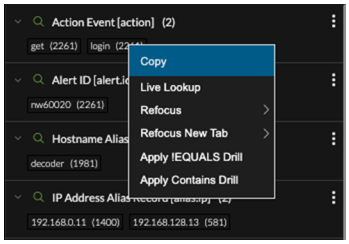
1. In the meta key row of an entry, click the Meta Key options button (⋮).
   The Meta Key options are displayed. Currently the only option is Copy Values.



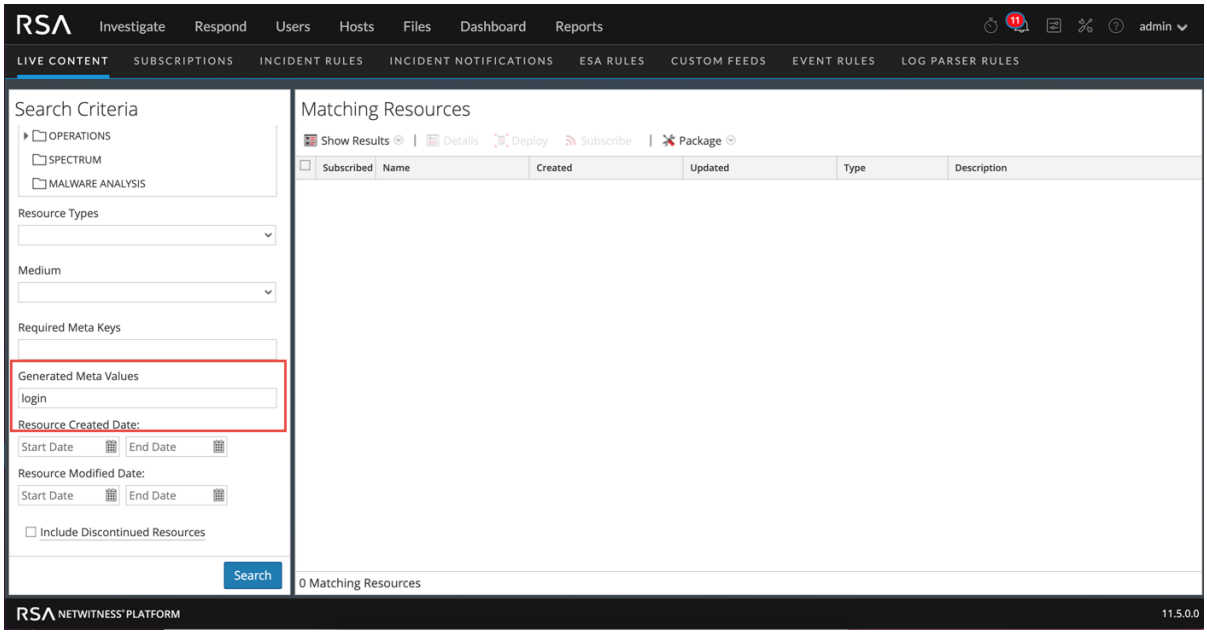
2. Click **Copy Values**.
   A comma-separated list of the values is copied to your local clipboard. This is an example of the clipboard contents: `"get", "login", "put"`.

View a Selected Meta Value in RSA Live

1. Right-click a meta value, for example **login**.
   The Meta Value drop-down menu is displayed with the Copy option selected initially.



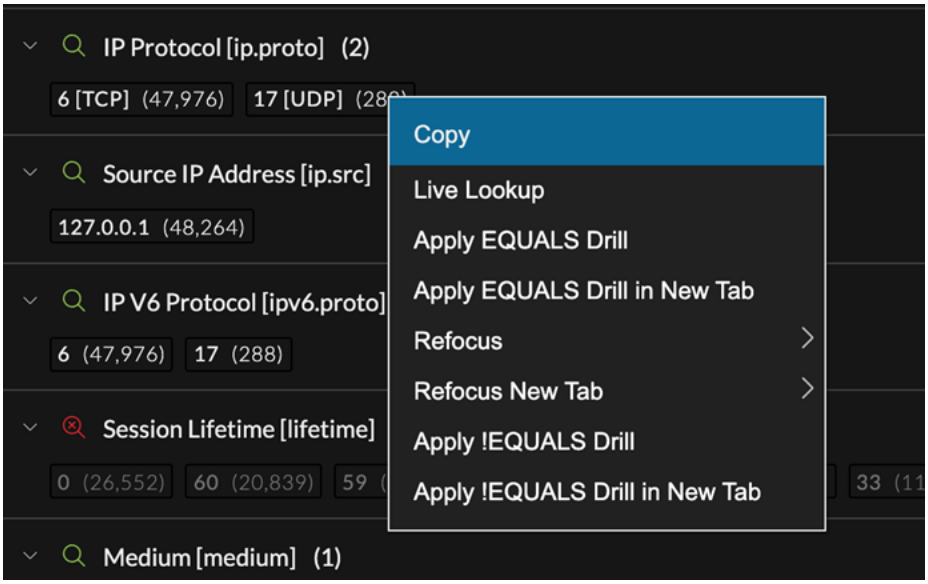
2. To look up the meta value, for example **success**, in RSA Live, select **Live Lookup**.
   The Live Search view is displayed with the meta value entered in the Generated Meta Values field, and ready for a search.



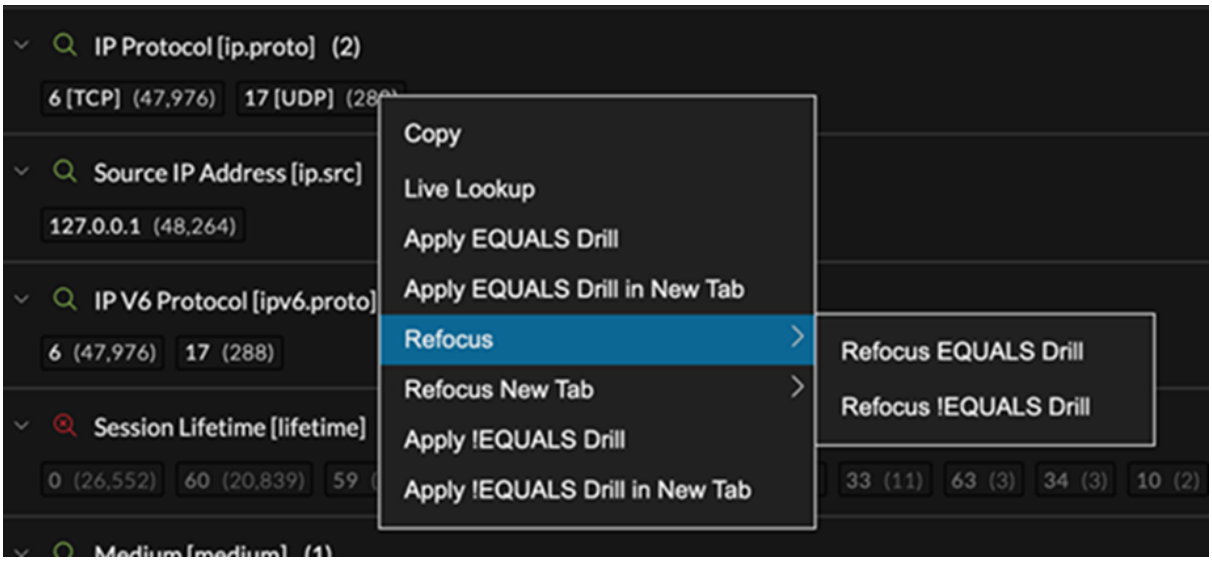
Refocus the Investigation of a Meta Value

For each value listed under a meta key, the focus is `<meta key> = <meta value>`. When you right-click a meta value, a context menu with different refocus options is displayed. All of the refocus actions update the drill point in the Events panel and the Filter Events panel.

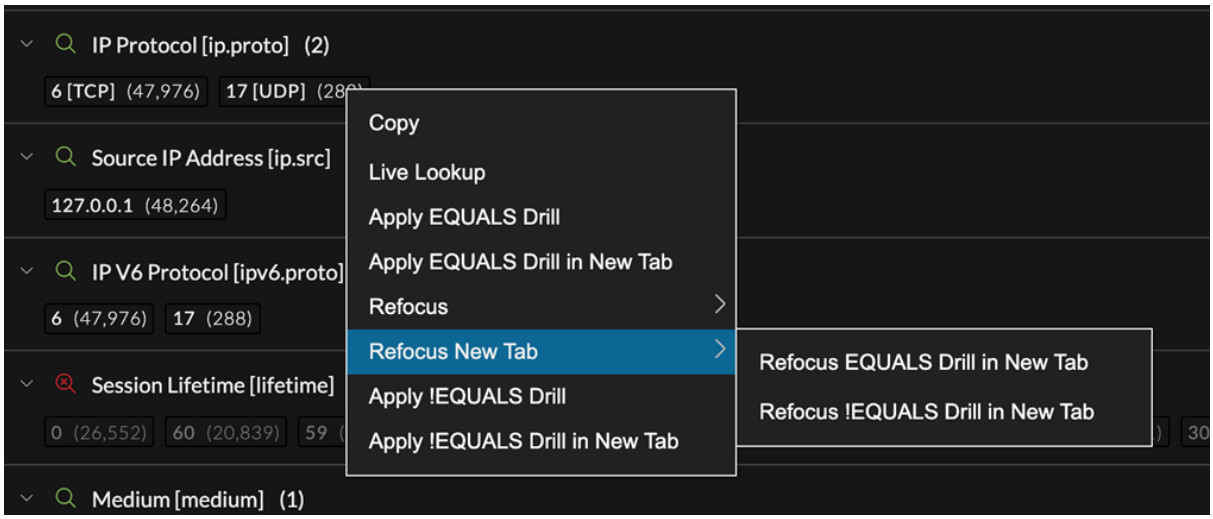
1. To append the key-value pair to the query with different operators (=, `!=`, `contains`), right-click a meta value (for example **UDP** in the figure below) and select one of the **Apply <operator> Drill** options.



2. To start the query over with the key-value pair and a different operator (=, `!=`, `contains`), right-click a value and select one of the **Refocus <operator> Drill** options.



3. To append the key-value pair to the query or start the key-value pair over in a new browser tab, right-click a value and select one of the **Refocus New Tab > Refocus <operator> Drill in New Tab** or **Refocus <operator> Drill in New Tab** options.



The drill is refocused according to your choice, and the new query is executed in the Events panel.

Filter Results in the Events View

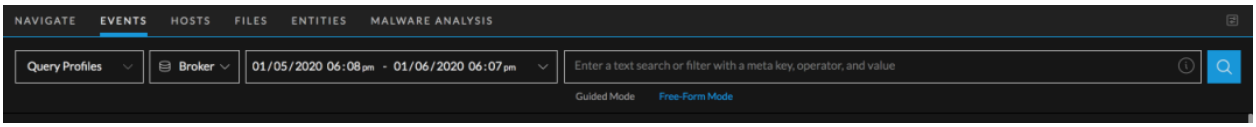
Note: This section applies to version 11.1 and later.

Filtering events in the Events view helps to narrow the focus of an investigation to a smaller, relevant set of events. You can filter events in the Events view using the Events Filter panel (Version 11.5 and later), the options in the query bar, and the options in the Events panel.

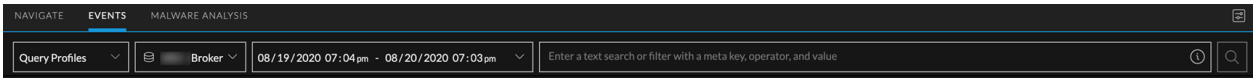
Initial Filter Using the Query Bar

When initially opening the Events view, the most basic filtering is to select a service, and time range and then query the service in the query bar. This returns a list of matching events in the Events panel. You can also select a query profile (Version 11.4 and later) and build a query to look for events that contain certain meta keys, meta values, and text in the query bar.

This figure illustrates the Version 11.4 and earlier query bar with the options to select a query profile, service, and time range to filter events and load them in the Events panel. Two modes are available Guided Mode and Free-Form Mode.

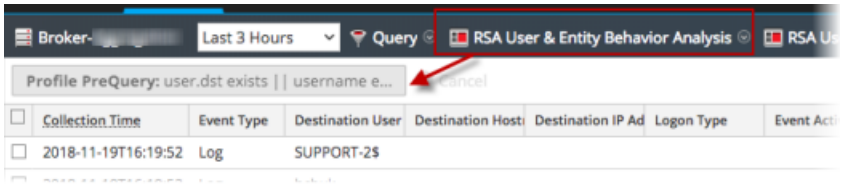


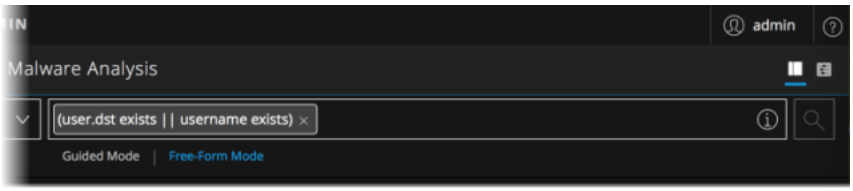
This figure illustrates the Version 11.4.1 and later query bar in which the Guided and Free-From Modes are no longer necessary. The simplified filter entry form allows you to use advanced auto-suggest options and also to enter a free-form query.



- The Query Profiles menu is available in Version 11.4 and later. You can encapsulate a query and a column group in a profile so that a useful combination of attributes is easily recalled and applied to a set of events in the Events panel (see Use Query Profiles to Encapsulate Common Areas for Investigation.)
- By default, the first service is automatically selected (unless you previously selected a service and the selected service is in browser cache). You can select a service as described in Begin an Investigation in the Events View
- If you do not select a time range, the default time range (24 hours) is used.
- The query builder field is an empty field to the right of the time range selector. This is where you build a query by creating filters. Clicking submits the query and sends a request to the selected service to load the data. In Version 11.3 and later, clicking the (console icon) opens the query console, where detailed status of the query is provided.
- When you go to the Events view from the Legacy Events view or the Navigate view, the service, time range, and any filters that were selected in the Legacy Events view or Navigate view are displayed in the query bar. The service, time range, and individual filters can be modified.

- If a profile is selected in the Legacy Events view when you right-click or double-click an event and go to the Events view, the filters from the profile (preQuery) are added to the query builder field as an editable filter. The following figures show a preQuery in the Legacy Events view, and the same query added as the first filter in the Events view.

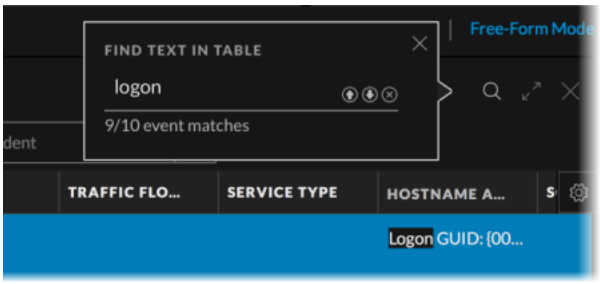




Find a Text String in the Events Panel (Version 11.4 and Later)

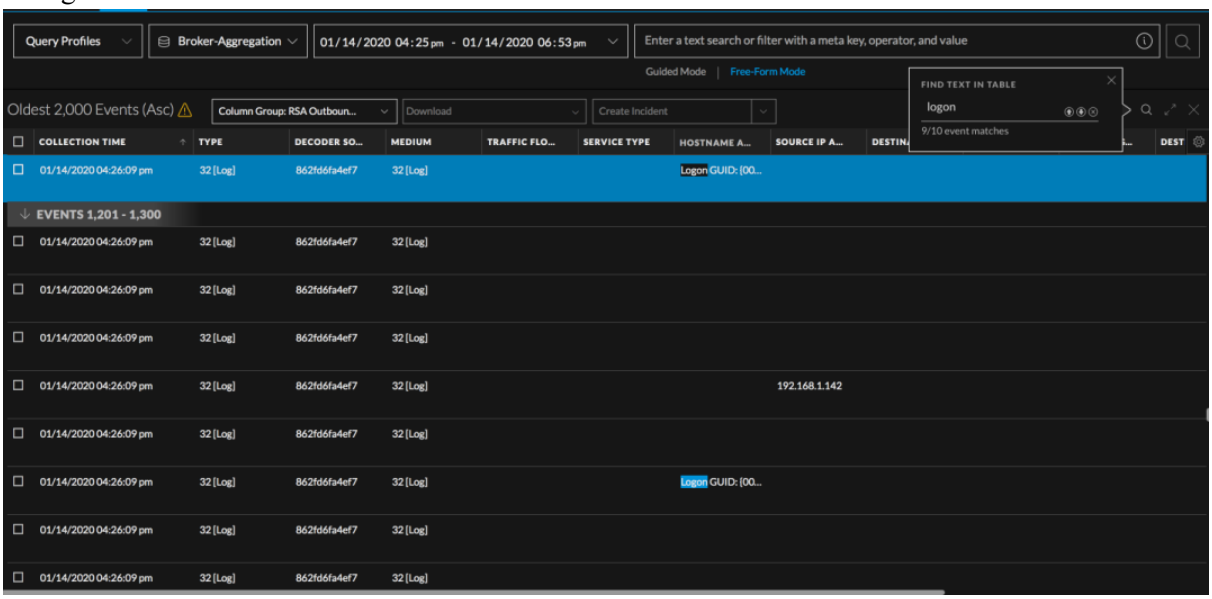
With the Events panel open, you can search for a text string in the list of events. This search is similar to the CTRL-F search in a browser window. The search scans all text in every row of the table, visible columns only, to find matching text and highlights the matches. Columns that are not displayed are not searched. The search function is disabled if the Summary column is part of the table.

1. With events loaded in the Events panel, click  on the right side of the toolbar.



2. In the **Find Text in Table** dialog, start typing a text string.
   After you type two characters, exact matches of the text string without regard to case are highlighted in the Events panel. As you type more text, highlighted events are further refined. The following figure is an example of the results found after entering "192.168" in the Find Text in Table dialog. The text string was found in 10 events. The first event is highlighted in blue with the text string within the event also highlighted. Icons are available for navigating the search results and closing the

dialog.



3. To navigate through the search results, click the up and down arrows.
   - To view the next event that contains the text string and navigate downward through the search results, click the down arrow. If you click the down arrow when viewing the last result, the first result is highlighted.
   - To view the immediately prior event that contains the text string and navigate upward through the search results, click the up arrow. If you click the up arrow while viewing the first result, the last result is highlighted.
4. To close the search dialog, click X or press the ESCAPE key. The dialog also closes if you open a reconstruction, select a new column group, or execute a new query.

Refining the Results in the Events Panel

After the initial filter and query submission, you can continue to use the options in the query bar to refine results, with two added methods of filtering the results.

- In Version 11.4, you can use column groups to optimize the number or attributes (meta keys, meta groups, meta entities) you look at for a given event (see Use Columns and Column Groups in the Events List).
- In Version 11.5 and later, you can filter events by exploring meta keys and meta values in the results in the Filter Events panel, which is a Beta release feature. This allows you to pivot through the metadata as you can in the Navigate view, and offers the added convenience of immediately seeing the matching events in sequence in the Events panel based on your drill point. The administrator can enable or disable this feature as describe in the *System Configuration Guide*.

Query Builder Concepts

In the query builder, you can reduce the number of events to an interesting set by creating three types of filters: simple, free-form, or text.

The basic syntax for each filter is as follows: `<meta key><operator><meta value>`. Here is an example: `direction = 'outbound'`.

In Version 11.4, when you type or paste a query in the query bar, the text is parsed into individual filters separated by the AND operator if the parsing engine determines that AND is needed. Earlier versions use only the AND operator between filters, and the logical operator is not visible.

- If you type `action = 'get' action = 'put'`, the result is two filters separated by AND.
- If you type `action = 'get' OR action = 'put'`, the result is two filters separated by OR.

When typing or pasting a filter for `event.time`, use one of the following formats:

- `event.time = '2020-DEC-02 23:00:00'`
- `event.time = '2000-12-20 21:00:00.000'`
- `event.time = '2000-12-20 21:00:00'`

In Version 11.4, the parsing engine converts a longer string of text that you type or paste in the query bar into individual filters. Parts of the filter that are not parsable are converted to a free-form filter. In earlier versions, a long text string is added to the query bar as a single filter. Further enhancement in Version 11.4.1 provides the ability to keep typing text for any query, where you type a meta key and an operator or an operator and a value, as a free-from query. The free-from query is parsed as usual.

- If you type, `action = 'GET' OR action is 20 || action = 'PUT'` in the query bar the Free-Form option is used. Part of this text cannot be parsed so the result is three filters separated by OR. The following figures show the query bars of version 11.4 and version 11.6 respectively.





- In Version 11.4.1, If you type a meta key-operator-value sequence and you continue typing without pressing Enter, the Free-Form option is automatically used so that you can continue typing the query. For example, you can type `medium = 1 OR medium = 2` without pressing the Enter key before `OR`. The Free-Form option is highlighted while you type and when you press Enter at the end, a free-form filter is created in the query bar.
- Text filters (Version 11.4 and later) are text strings that do not contain spaces. You can search the data set for any exact match of indexed meta keys, not all meta keys. Here are some examples: `failed`, `login`, or `attempt`.

Note: In some cases, when you are typing a text filter that is close to matching a meta key and operator statement, the auto-suggest feature erroneously suggests a filter using the meta key and operator. The workaround is to begin typing the text and select Text Filter at the point where the auto-suggest feature turns the text into a meta key and operator. For example, there is meta key named `crypto`, and an operator named `contains`, and you want to create a text filter to search for `cryptocurrency`. As you type `c-r-y-p-t-o`, the next `c` in "currency" triggers the `contains` operator instead of continuing to type as a single word. To complete the text filter, right before typing that `c` in `currency`, which would trigger the `contain` operator, highlight the Text Filter option to let the system read the input as text filter.

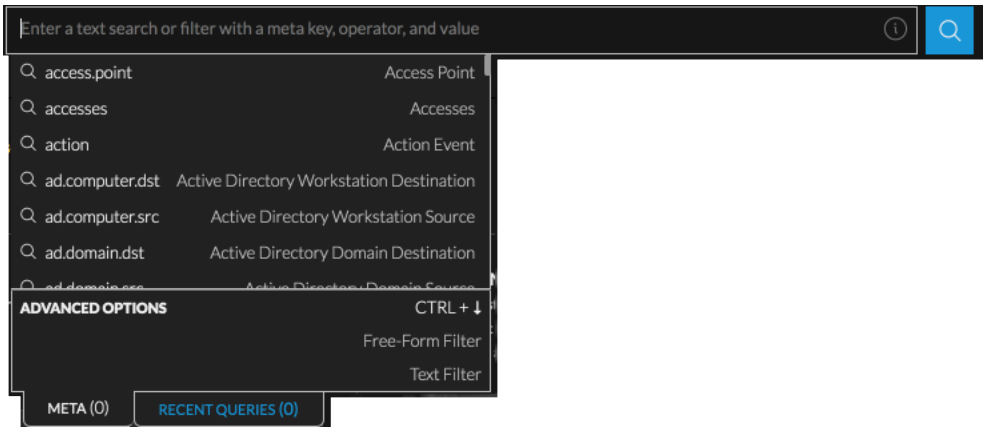
In the query builder, each filter becomes an editable field. Filters line up from left to right, representing the sequence in which the filters were created. As more filters are added and exceed the length of a single line, they wrap to another line and the input area expands vertically so that all filters are visible without scrolling to the right.

Successive versions of RSA NetWitness® Platform 11 have more capability than the original 11.1 query bar, with Version 11.4 offering extensive help as you create a query.

Guided Mode vs. Free-Form Mode

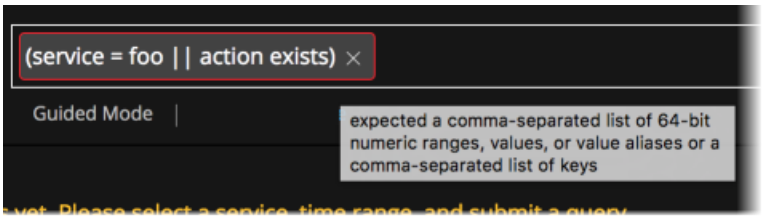
Note: Version 11.4 offered two modes for entering queries in the filter entry form: Guided Mode and Free-Form Mode. Beginning with Version 11.4.1, the powerful auto-completion features and suggested values of Guided Mode and the ability to type or paste a free-form query are fully integrated. References in this document that differentiate between Guided Mode and Free-Form Mode are for analysts using Version 11.4.0.x and earlier.

In Guided Mode, you are guided with suggestions for auto-completion that show valid meta keys and operators, and suggested values in the filter entry form. In version 11.4, you can type, paste, choose a recent query, or select from the drop-down menu. Earlier versions do not support pasting text and recent queries. This is an example of the 11.4 filter entry form.



As you create filters, the syntax of each filter is validated and invalid filters are marked by a red outline. If you hover the mouse over the filter, a message that explains the error is displayed.

In Version 11.3 and later, free-form filters are validated on the server side, which may take additional time. If you submit the query before the server has returned filter validation results, the is replaced by a spinner. When server validation returns, a query with no invalid filters begins execution. If the query contains an invalid filter, execution is terminated and the invalid filter is outlined in red. This is an example of an invalid query.



In Free-Form Mode, you can type or paste a long text string. There is no auto-suggestion, and validation is performed on the server side when you submit the query. If an error is found, the query does not execute.

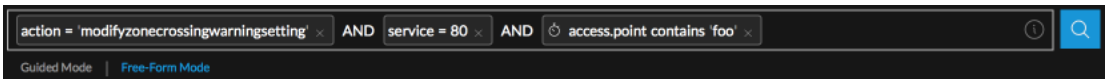
Note: The button has a different label in versions earlier than Version 11.3. It was previously named **Query Events**.

Clicking Guided Mode or Free-From Mode toggles between modes. If you selected Free-Form Mode the last time you logged in, this choice is stored in browser cache and is used until the browser cache is cleared.

- When you switch from Guided Mode to Free-Form Mode, filters that you created in Guided Mode are transformed to a text query in the Free-Form field.
- When you switch from Free-Form Mode to Guided Mode, the query you were typing is added to the query bar as individual simple filters, but it does not include auto-suggest options.

  Note: Before Version 11.3, a Free-Form filter could not be edited in Guided Mode.

The following figure is an example of the query bar with the Guided Mode query builder with several filters.

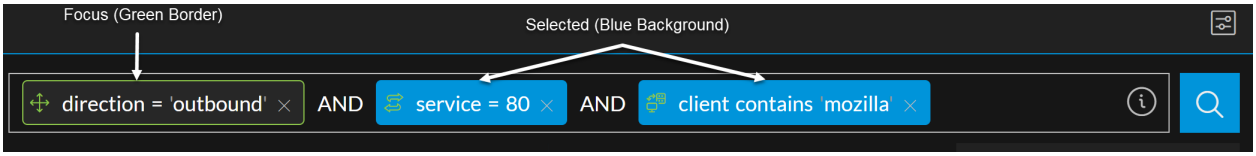


The following figure is an example of the Free-Form query builder in use.



Concepts for Editing Multiple Filters

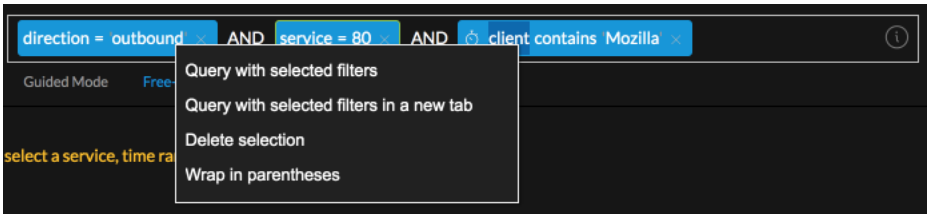
As you work in the query builder, you can see when a filter has focus for editing (a green outline) and which filters are selected (blue background). This is useful because you can have multiple filters selected for right-click actions, but only one can be edited at a time. The figure below shows the green outline marking a filter that has focus and the blue background indicating that two filters are selected.



This figure illustrates the same set of filters with all filters selected (blue background) and one filter that has focus (blue background and green outline).

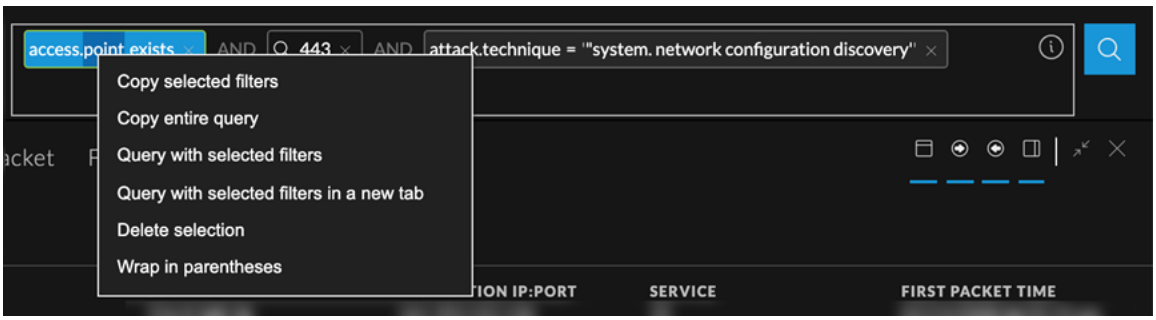


A right-click action from the drop-down menu applies to all selected filters as shown in this figure showing Version 11.4 options.



In Version 11.4.1, the menu has two new copy options as shown in the following figure. The options allow you to share the clipboard contents with other analysts or paste the contents to the query bar. You can:

- Select a single filter, right-click it, and then copy the entire query to the local clipboard.
- Select multiple filters, right-click one of them, then copy the selected filters.



These are a few basic concepts that explain how to work in the query builder:

- You can select multiple filters, but only one can have focus and the last selected filter is the one with active focus at any point in time.
- To select a filter and give it focus, click the filter. To deselect the filter and remove focus, click the filter again, press **Esc**, or click anywhere else on the page.
- To add a filter, click before or after an existing filter. To create a new filter before or after the filter in focus, press the right or left arrow key.
- To open a filter for editing, double-click the filter or click it and press **Enter**. To exit without saving changes and leave the filter in focus, press **Esc**.
- To delete a filter, click the filter and press **Delete** or click **X** on the filter.
- (Version 11.6) If a filter is awaiting user's input, it is highlighted with blue color .



- (Version 11.6) An invalid filter is highlighted with red color.



- You can hover over the filters in the query bar to display tooltip messages.

The Version 11.3 and Earlier Query Builder

In Version 11.1, the user interface guides you as you create and edit each simple filter: `<meta key> <operator> <meta value>` filter. The user interface supports only simple filters. If you open an event from the Legacy Events view or Navigate view and the filter has more than one operator, `||`,`&&`, `()`, `REGEX`, or `LENGTH`, the filter is added, but editing is not supported in the Events view. Refer to the Investigate User Guide for NetWitness Platform 11.3 for details; it is available here in PDF form: Master Table of Contents for RSA NetWitness Platform Version 11.

In Version 11.2, the user interface has two modes: Guided Mode and Free-Form Mode. Guided Mode guides you as you create and edit each simple `<meta key> <operator> <meta value>` filter. The default mode is Guided Mode, which includes auto-suggest and validation options. You can enter a long text string in Free-Form Mode. In Free-Form Mode validation is performed when you run the query. Refer to the Investigate User Guide for NetWitness Platform 11.2 for details; it is available here in PDF form: Master Table of Contents for RSA NetWitness Platform Version 11.

In Version 11.3, the user interface has these added features:

- When you go to the Events view from the Navigate or Events view, with a preQuery in effect for a profile, the preQuery that was displayed in the breadcrumb is available in the Events view query builder as an editable filter.
- The auto-suggest options in the user interface are augmented with an Advanced Options section that allows for the creation of a free-form filter. Free-Form Mode is still useful for pasting an entire long text string.
- You can cancel a query while it is executing.
- You can view detailed status information about the query being executed in the query console.

The Version 11.4 Query Builder

You can type; select meta keys, operators, and values from the drop-down menus; or paste a filter in the query bar. Added 11.4 features in the Guided Mode filter entry form are described in detail below.

Meta Keys Cached for Faster Loading

When the Events view opens, meta keys from all connected services are cached for faster data loading. These meta keys are available in user interface elements that have auto-suggest meta keys. (When you are building a column group or profile, if you are expecting to see a meta key and it is not displayed, select the service where the key was added to force a cache update. This usually occurs only when a meta key is not added to all concentrators.)

Text Filter

You can create a text filter to find a text string in the data set which is indicated with  . You can use a text filter with no knowledge of meta keys that would contain the values. One text filter per query is supported. A text filter looks through indexed meta keys, not all meta keys.

Pasting Text Instead of Typing

When creating a filter, you can paste a meta key or value in the filter entry form. When you paste text into the filter entry form instead of typing the text, the text is parsed appropriately to create one or more filters. Any portion of the text that cannot be parsed is converted to a free-form filter.

Select All Filters and Copy All Filters (Version 11.4.1)

While creating a filter in the Events view query bar, you can use keyboard commands to select all filters (Cmd-A for MacOS, Ctrl-A for Windows) and then copy the selection to the clipboard (Cmd-C for MacOS, Ctrl-C for Windows). The clipboard text is available to share with other analysts or to paste in the query bar using Cmd-V or Ctrl-V).

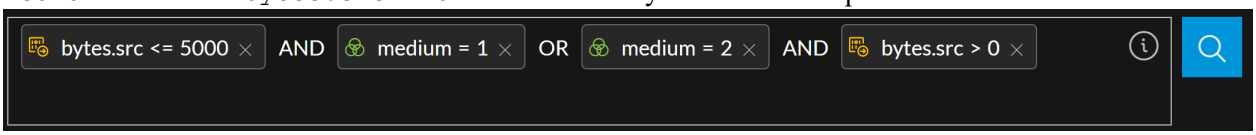
Use of Recent Queries

The filter entry form offers two methods of entering meta keys, operators, and values: the Meta tab and Recent Queries tab. The Meta tab is the same as the filter entry form for prior versions except that a count of matching results is given in the tab label and icons mark meta keys that are indexed by key, indexed by value, and not indexed. In the Recent Queries tab, up to 100 recent queries are displayed. The list is filtered as you type to show only queries that contain the typed text, and you can select a query from the list.

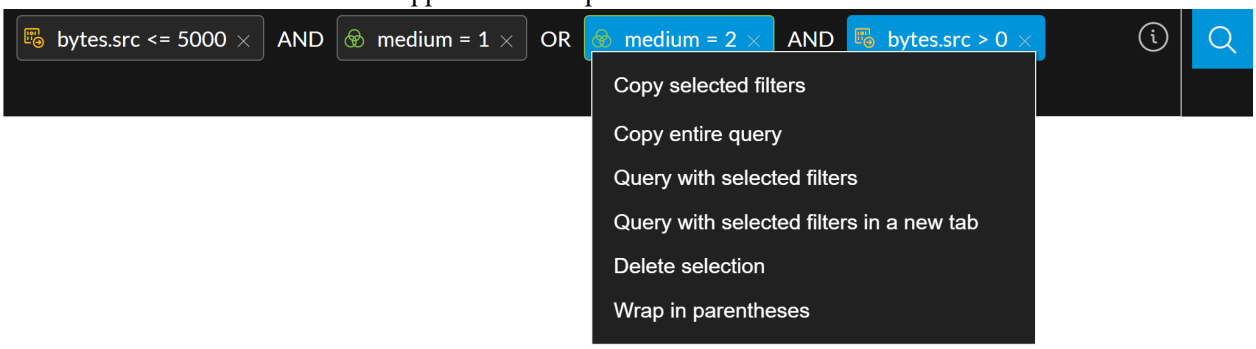
Use of Advanced Operators

Auto-suggest can parse the following advanced operators that you paste or type into the filter entry form: `<`, `>`, `<=`, `>=`, `OR`, `||`, `AND`, `&&`, `()`, `regex`, and `length`. The text is parsed as multiple filters. For example, if you type or paste `medium > 0 && medium <= 100`, the text is parsed as two simple filters with an explicit AND operator: `medium > 0 AND medium <= 100`. If you type or paste `bytes.src <= 5000 && medium = 1 || medium = 2 && bytes.src > 0`, four simple filters are created with `AND` and `OR` operators separating them: `bytes.src <= 5000 AND medium = 1 OR`

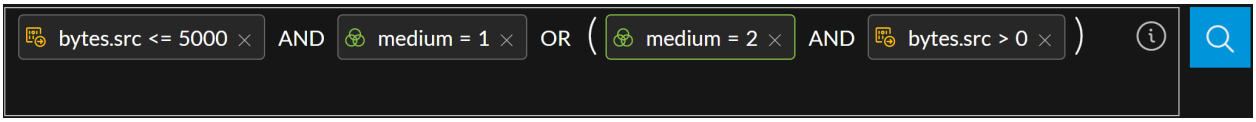
`medium = 2 AND bytes.src > 0` to make as many valid filters as possible.



This filter is an example of a filter in which it would be useful to add parentheses. You can select `medium = 2` and `bytes.src > 0`, then right-click and select **Wrap in Parentheses** from the drop-down menu. Text filters are not supported inside parentheses.



The resulting query is `bytes.src <= 5000 AND medium = 1 OR (medium = 2 AND bytes.src > 0)`.



If you encounter errors while creating filters, look for tooltip messages and check the documentation.

Easy Use of AND/OR Operators

When you type `||` and `&&`, they are displayed as `OR` and `AND` in the query bar. You can change `OR` to `AND` and `AND` to `OR` by clicking the word. When you insert the cursor to add a filter, the `AND` operator is added before the cursor. When you delete a filter, orphaned `OR` and `AND` operators are removed. The operator for a text filter must be AND because text filters are always ANDed to a query.

Automatically Balanced Parentheses

When creating and editing filters in the query builder, parentheses pairs are automatically balanced as you type. If you type an open parenthesis in a filter that is open for editing or before a selected filter, the close parenthesis is added at the end of the filter. This works intuitively as you type so that you can add new filters on either side of the parenthesis and between parentheses when there are nested parentheses. Orphaned parentheses are automatically removed. If adding parentheses would create an invalid filter, the parentheses are not added. You can also right-click selected filters, and add parentheses using the Wrap in parentheses option. This option is only available when the result would be a valid filter.

Hints about Available Values

For properly indexed meta keys, the user interface provides hints about available values related to the time range of the query. Up to 100 suggested values are returned and, when you type text, the list of 100 values is filtered to include only relevant values. If no matching values are returned, a message advises "No suggestions found." (The suggested values are based solely on the time range; filters in the query do not filter the list of 100.)

CIDR Notation and Shorthand

When entering a value for an IP address in a filter, you can use CIDR notation to filter for addresses within a range.

The IPv4 CIDR block range is 0 to 32. For example `10.20.30.0/24` specifies `10.20.30.0` with a subnet mask of `255.255.255.0`, which will match an IP in the range `10.20.30.0` through `10.20.30.255`.

The IPv6 CIDR block range is 0 to 128, for example,
`1203.0fe1:fe82:b896:89b0:8a7c:99bf:323d/32` specifies
`1203:0fe1:0000:0000:0000:0000:0000:0000` through
`1203:0fe1:ffff:ffff:ffff:ffff:ffff:ffff`.

You can also use shorthand to remove groups of zeros or leading zeros in a group in IPv6 addresses, for example,

`1203:fe1::`

There must be no spaces between the IP address and the CIDR mask that you are using.

Ranges or Series of Values

For meta keys that have numerical data, you can use a range of values, a series of values, or both to filter data. For example, this query has a comma-separated list, and two of the values in the series are ranges `src.port = 0-1023, 1024-1050, 65535`. If a comma is a part of a value, the value must be wrapped in quotes. For example, `get,post` is interpreted as two separate values, while `'get,post'` is interpreted as one value. A range of values must be a valid range of positive integers, separated by a dash (with our without a space before and after). The first number in the range must be smaller than the second. For example, `0-1023` and `0 - 1023` are valid ranges, but these are not valid ranges: `-10 - 50`, `50 - 10`, `50.8 - 60.2`, `50 - 70x`.

No Separating Space Required After Meta Key and Operator (Version 11.4.1)

Filters in the query bar need a space between the meta key and the operator, and between the operator and the value. Operators must be typed with a separating space in the filter entry form in order to use the auto-suggest functions for operators and values. To improve the user experience when typing a query, the filter entry form accepts operators typed with no separating space after the meta key. When you type an operator with no separating space, a value is auto-suggested as usual and a space is added between the meta key and the operator. When you type an operator and a value without a separating space, the space between them is automatically added.

Select a Time Range

The Time Range selector limits the events returned in the Events view to a specific time range. The time range is displayed in the format `Start Time - End Time`, showing the date, hours, and minutes in your current timezone based on timezone settings configured for your profile. In Version 11.3 and later, you can choose a time range relative to the current collection time or create a custom time range. The time and date format is based on preferences set for the Events view in the User Preferences dialog (select  > **Profile**).

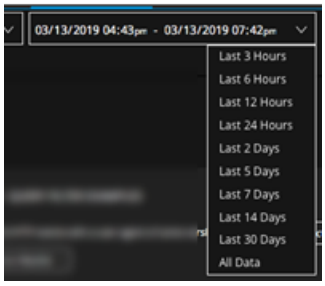
- By default, the date format is MM/DD/YYYY. You can change the format to DD/MM/YYYY, or YYYY/MM/DD in the User Preferences dialog.
- Start time and end time are in the format of HH:MM. Although seconds are not displayed, the value for start time always defaults to HH:MM:00 seconds, and the value for end time always defaults to HH:MM:59 seconds. As an example, a time range of 6:45 pm - 7:45 pm is interpreted as 06:45:00 - 07:45:59 pm.
- The default time range is the 24-hour clock; you can change it to 12-hour periods.

Note: By default, the time format for downloads is Epoch format, which shows the time as a numerical value representing the number of seconds from the Unix epoch, January 1, 1970. The resulting number requires a conversion to be understood. Your administrator can change the setting for time format in downloads to combine your user preference time zone, date format, and time format into an easily understood representation, which follows the industry standard ISO 8601 representation when possible. Given this time: `04/13/2020 09:17:36 am` with timezone US/Pacific `(GMT-7:00`, this is an example of the time on the 12-hour clock as it appears in the user interface: `04/13/2020 09:17:36 am`. In the download, Epoch format would represent the time as `61547519856000`. If your administrator set the time format for downloads to the easily readable representation, the same time would be represented as follows: `04-13-2020T09:17:36AM-07:00`.

The time format for a query is based on preferences set for the Events view in the Event Preferences dialog (select). The time format can be either database time or wall clock time. When Database Time is selected, the start and end time for a query is based on the time that the event was captured (collection time). When Wall Clock Time is selected, the query is executed using the end time based on the current browser time; the start time is calculated based on that end time and the time range. This and other Events view preferences are described in Configure the Events View.

To edit the time range, do one of the following:

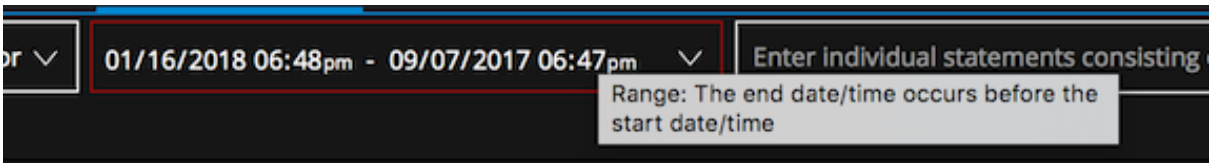
1. Click the drop-down arrow inside the **Time Range** selector and select a time range from the list. Options are in minutes, hours, days, or all data.



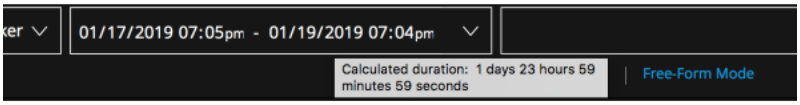
2. (Version 11.3 and later) Edit the time range directly by clicking the year, month, day, hour, or minute displayed in the query bar. When a value is highlighted, type a new value for either the start or end time. If your time format preferences are set to 12-hour periods, click **am** or **pm** to toggle between the two options.



If the time range is invalid (for example, the start time is later than the end time), a red border appears around the Time Range selector. The button is disabled because the query is no longer possible, and a tool tip shows an error message explaining what you need to change. The following figure shows an invalid time range.



The selected time range is stored in your browser for the service being queried; you can set different time ranges for different services. A tool tip shows the calculated duration of the query. The following figure is an example of the tool tip.



Submit a Query

The button on the right side of the query bar is active as needed to submit a query. In version 11.3 and earlier, when you click , all of the filters are ANDed to generate results and the button becomes inactive. In Version 11.4, because the query may contain other operators besides AND, the query is submitted as is. The button becomes active again in these conditions:

- If you change the service in the query bar or change the column group in the Events panel, a network call for data for a reconstruction in the Events panel continues to use the previous service, time range, and metadata filters until you submit the new query. The button becomes active as an indicator that the data in the view is stale.
- If more than a minute has passed and the original query's time range would no longer generate the same result set, the button becomes active as an indicator that results may be stale. In Version 11.3 and later, a setting in the Events view preferences determines this behavior by enabling or disabling the Update Relative Time Window Automatically option (see Configure the Events View.)

Cancel Execution of a Query

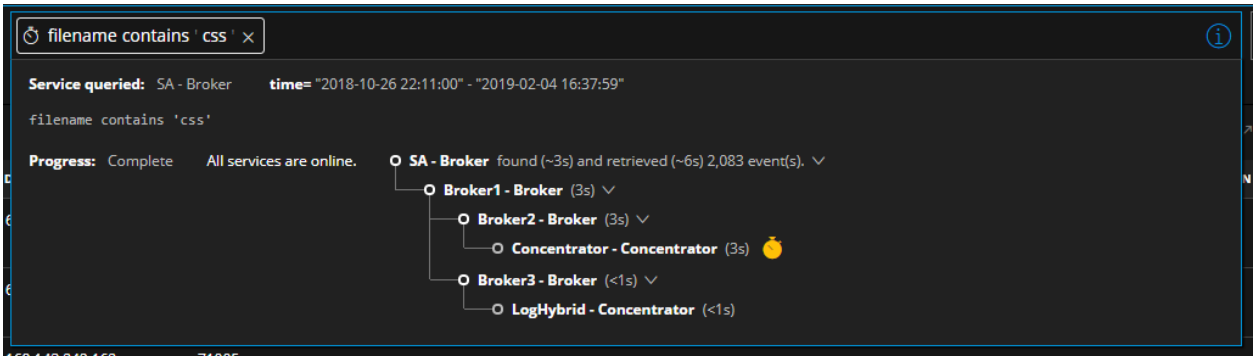
After you click to submit a query, the button changes to (the stop query option). The stop query option remains until all the events are loaded in the Events panel. To cancel the query, click .

If the query is canceled before all results have been returned, the following message is displayed at the end of results in the Events list: "Because the query was canceled, only partial results are displayed."

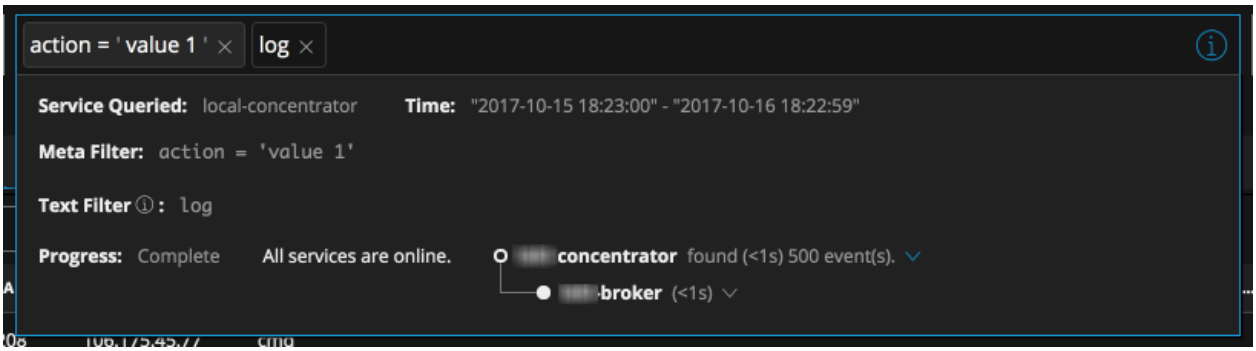
View Status of a Query

After a query is submitted, you can click the information icon ( ) in the query bar to open the query console. In the query console, you can see which service, time range, and metadata was queried as well as real-time information about the status of the query and the services being queried. The time range displayed in the Query console always shows each date as YYYY-MM-DD. Here is an example range from the Query Console: `"2014-09-20 20:57:00"-"2018-11-02 18:57:59"`.)

The following figure is an example of the Query Console for Version 11.3 when a query executes successfully and the slowest service is marked by an amber stopwatch.



The following figure is an example of the information in the Query Console for Version 11.4 after a query that includes a text filter executes. Notice that the query is shown in two fields, Meta Filter and Text Filter.



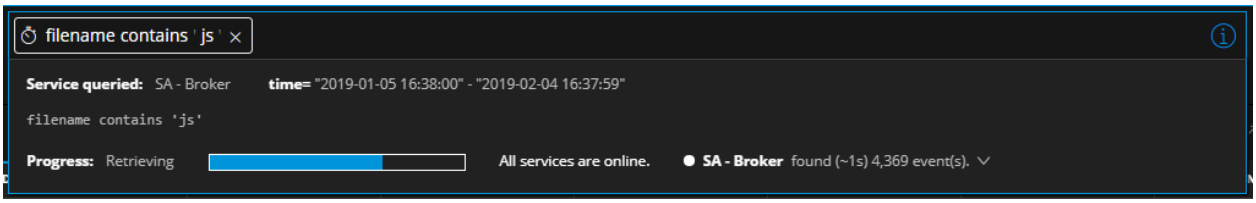
While a query is executing, a progress bar indicates the query's completion percentage at the bottom of the console. The status lets you know details about what is happening; for example, you can tell when the query is executing, queued, reading the index file for the queried service, retrieving events, and complete. All statuses and non-fatal messages are displayed as they come in, and the border color of the query bar changes to amber if a non-fatal error occurs.

Icons provide additional information about individual services.

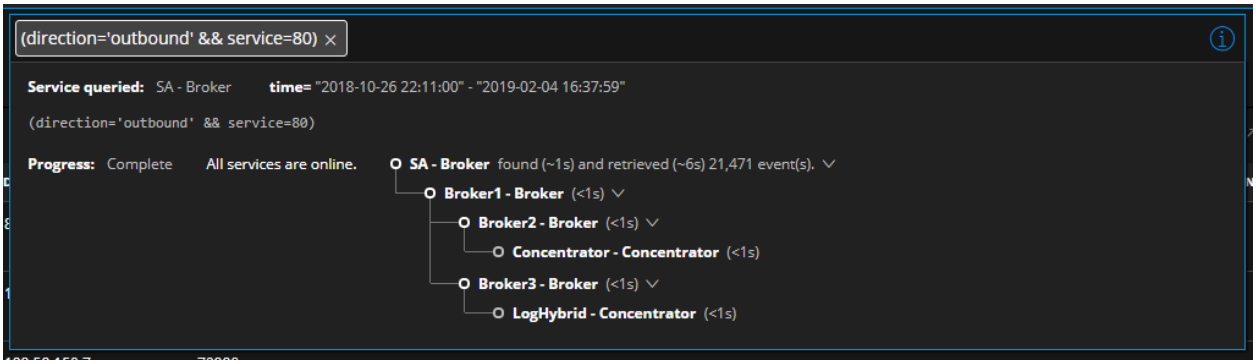
- An amber stopwatch marks the slowest service.
- An amber triangle indicates a warning was received.
- A red triangle shows that an error was received when trying to query the service.

Executing and Reading the Index File to Find Events. The first stage of a query is complete when the queried services have found results. The query console provides a nested hierarchical listing of all the services being queried with indicators showing which are online or offline, and the time in seconds that the service took to find results.

Retrieving Events and Loading in the Events Panel. While the found events are being retrieved and loading in the Events panel, the progress bar shows a visual indicator and text description of what is happening. In the figure below, results were found and are being retrieved.

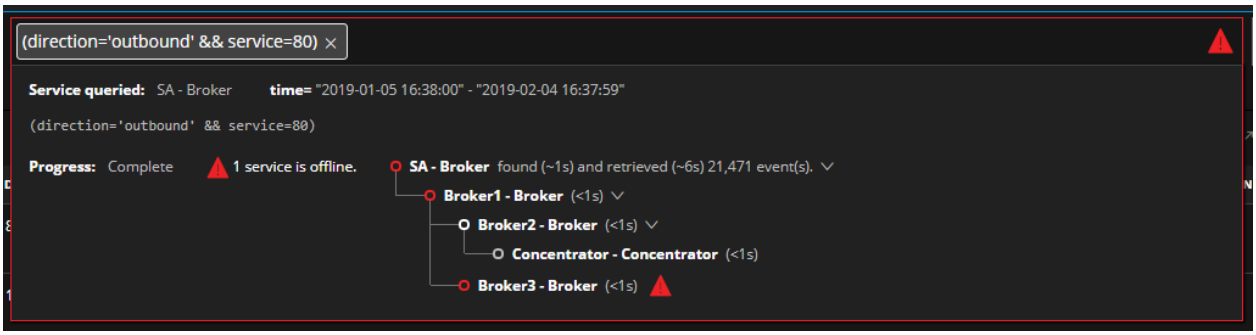


Request Complete. If there are no errors or warnings when loading is complete, the query console is outlined in blue and the button is disabled as an indicator that the data in the view is fresh. The following figure is an example of the query console for a completed query with no errors or warnings.



Errors and Warnings. A fatal error such as a syntax error in the query, or the queried service being offline, stops execution of the query. A red triangle is displayed in the upper right corner of the query console, and the console is outlined in red to indicate that the query failed. If the queried service is offline, only the queried service with no hierarchy of services is listed in the query console and marked by a red triangle.

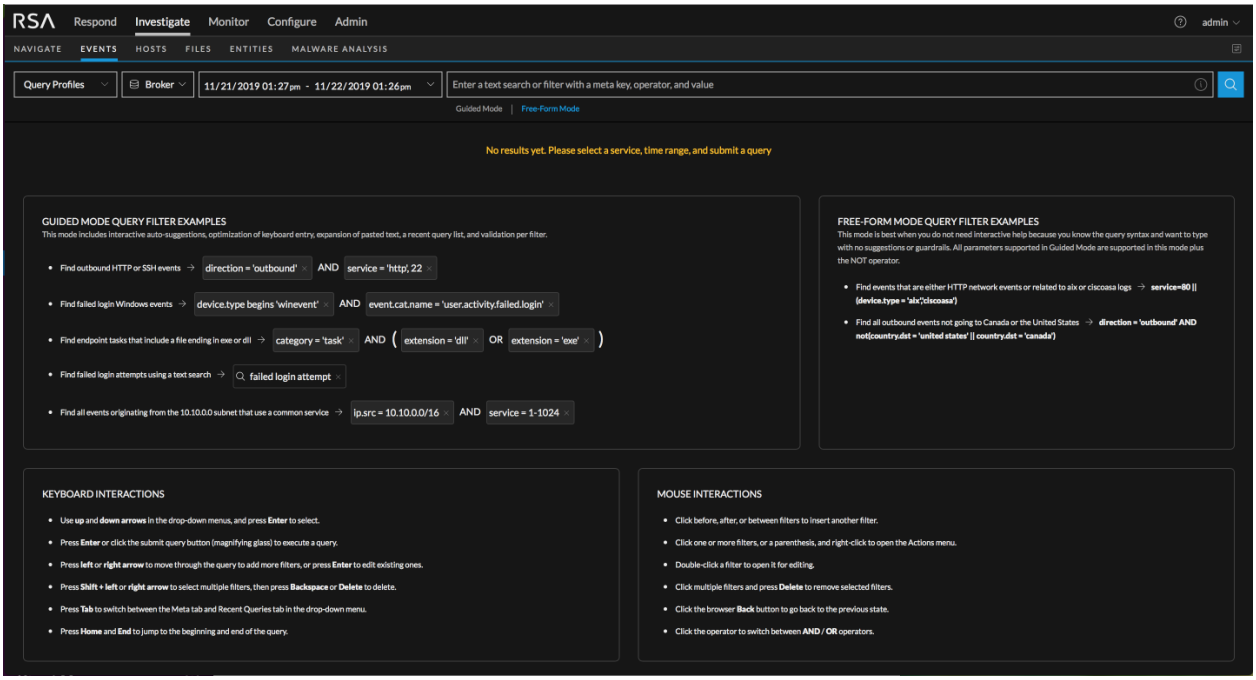
A non-fatal error does not prevent a query from executing. The query is executed and events are loaded, but a red triangle is displayed in the upper right corner of the query console, and the console is outlined in red as a warning. The following figure shows the appearance of the query console when the queried service proxies to another service that is offline.



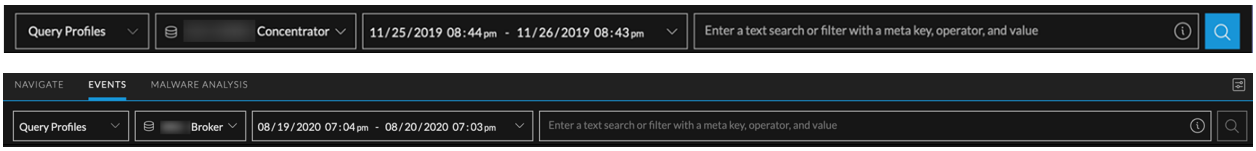
A warning does not prevent a query from executing. The query is executed and events are loaded, but an amber triangle is displayed in the upper right corner of the query console, and the console is outlined in amber.

Build a Query in Guided Mode

Guided Mode is the easiest way to create a query with features to help analysts enter valid queries. The following figure illustrates the initial Events view with Guided Mode in effect in the query bar (Version 11.3 and earlier).



There is no longer a need to select Guided Mode because all Guided Mode features and more improvements are built in to Version 11.4.1. This figure depicts the query bar in Version 11.4.1.



Keyboard Actions to Use in Guided Mode

In Guided Mode, the query builder allows entry, editing, and deletion of filters using the key strokes without having to use a pointer. Although you can use the pointer, you have the option to keep your fingers on the keyboard. This table identifies the available keyboard actions in Guided Mode when the cursor is located in the query bar; these do not apply to the service selector and time range.

Action	Keyboard Entry
Copy all filters (Version 11.4.1 and later)	With the cursor in the query bar, but not in a filter being edited, and all filters selected, press **Ctrl-C** (Windows OS) or **Cmd-C** (MacOS).

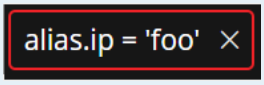
Action	Keyboard Entry
Delete characters in a filter	Selected characters: With characters selected in the query bar, press **Delete** or **Backspace**. Previous character (Version 11.4 and later): With the cursor next to a character in the query bar, press **Backspace** (Windows OS) or **Delete** (MacOS). All characters (Version 11.4 and later): With the cursor in a filter, press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS).
Delete filters	Selected filters: With one or more filters selected do one of the following: • **Right-click > Delete selected filters** or **Delete selection** (11.4 and later). • Press **Delete**. • Press **Backspace**. Filter that has focus (Version 11.4 and later): With the cursor in a filter that has focus, press **Backspace** (Windows OS) or **Delete** (MacOS). The focused filter is deleted and focus moves to the left. Filter that has focus (Version 11.4 and later): With the cursor in a filter that has focus, press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS). The focused filter is deleted and the focus moves to the right.
Delete parentheses in a filter, do not delete the contents (Version 11.4 and later)	With a set of parentheses, but not the contents selected, press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS). The selected parentheses are deleted, but the contents of the parentheses remain.
Delete parentheses and their contents in a filter (11.4 and later)	Selected parentheses: With a set of parentheses selected, do one of the following: • **Right-click > Delete selection**. • Press **Backspace** (Windows OS) or **Delete** (MacOS). The selected parentheses and contents are deleted and the focus moves to the left. • Press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS). The selected parentheses and contents are deleted and the focus moves to the right.
Deselect all filters	With a filter selected, press **Esc**.
Edit a selected filter	With a single filter selected, press **Enter**.
Insert a new filter at the beginning of the query bar, and open for editing (Version 11.4 and later)	With a filter selected, press **Home** (Windows OS) or **Fn** + **Left Arrow** (MacOS).

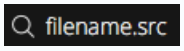
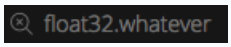
Action	Keyboard Entry
Insert a new filter at the end of the query bar, and open for editing (Version 11.4 and later)	With a filter selected, press **End** (Windows OS) or **Fn + Right Arrow** (MacOS).
Insert a new filter to the immediate left of the selected filter, and open for editing	With a filter selected, press **Shift + Left Arrow**.
Insert a new filter to the immediate right of the selected filter, and open for editing.	With a filter selected, press **Shift + Right Arrow**.
Insert a new filter to the immediate left of the selected filter	With a filter selected, press the **Left Arrow**.
Insert a new filter to the immediate right of the selected filter	With a filter selected, press the **Right Arrow**.
Open a new tab with the selected filters	With filters selected, **right-click > Query with selected filters in a new tab**.
Query with the selected filters	With filters selected, **right-click > Query with selected filters**.
Query with content of parentheses (Version 11.4 and later)	With parentheses selected: • To query with the selected parentheses contents, select one side of a parentheses set and **right-click > Query with selected filters**. • To query with the selected parentheses contents in a new browser tab, select one side of a parentheses set and **right-click > Query with selected filters in new tab**.
Select all filters in the query bar (Version 11.4.1 and later)	With the cursor in the query bar, but not in a filter being edited, press **Ctrl-A** (Windows OS) or **Cmd-A** (MacOS).
Select all filters to the left of the current filter	(Version 11.3.x and earlier) With a filter selected, press **Shift + Up Arrow**. (Version 11.4 and later) With a filter selected, press **Shift + Right Arrow** twice.
Select all filters to the right of the current filter	(Version 11.3.x and earlier) With a filter selected, press **Shift + Down Arrow**. (Version 11.4 and later) With a filter selected, press **Shift + Right Arrow** twice.
Select the filter to the immediate left if one exists	With no filter selected, press the **Left Arrow** key.

Action	Keyboard Entry
Select the filter to the immediate right if one exists	With no filter selected, press the **Right Arrow** key.
Submit a query.	With focus on the query bar and no pending filters, press **Enter**.

Visual Feedback in Guided Mode

Guided Mode provides visual feedback during query construction. This table identifies and describes the possible feedback.

Feedback	Icon	Description
Blue background on a Filter		Indicates that a filter is selected.
Green circle between two filters		(Version 11.3 and earlier) A green circle indicates the location of the cursor between two existing filters. Clicking inserts a new filter at this location. (Version 11.4) A bold cursor indicates the insertion point.
Green filter outline		Marks the single filter that has focus and ready to edit. This is combined with the blue background, when multiple filters are selected and this filter has focus.
Red filter outline		Indicates that the filter is invalid. A tool tip that explains the error is displayed.

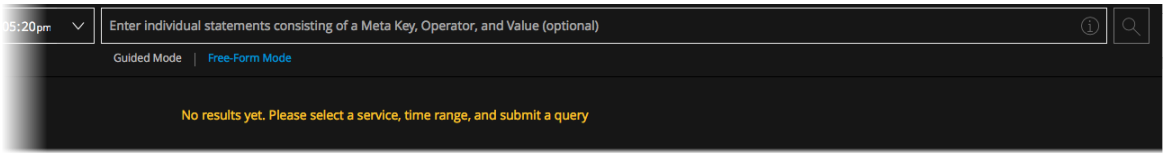
Feedback	Icon	Description
Index indicators in the Meta tab		(Version 11.4 and later) Indicate the index level of the meta keys in the Meta tab, which determines if you can use it in a filter:  This meta key is indexed by meta value and can be used in a filter.  This meta key is indexed by meta key, and can be used in a filter.  This meta key is not indexed, and not selectable for a filter. The `sessionID` meta key is a special case. Unlike other non-indexed meta keys, it is not configurable, but you can use it in a filter so it is marked by the key symbol. Supported operators are `exists`, `!exists`, `=`, and `!=`.
Query Events button		Used to submit a query, show the status of the query, and cancel a query. The button has three possible states:  Ready to submit a query using filters in the Query Builder.  Waiting for server validation to complete before executing the query.  The query is executing, click to cancel execution.
Slow Service icon		In the query console, marks the service that took the longest time to load results from the query.
Spinner in the Events list		Indicates that the query is currently being processed. The **Query Events** button is disabled while this occurs.

Feedback	Icon	Description
Stopwatch		(Version 11.5 or earlier) Indicates that the meta key/operator combination requires extra time to process. While the query is still executable, a more efficient meta key or operator is recommended.

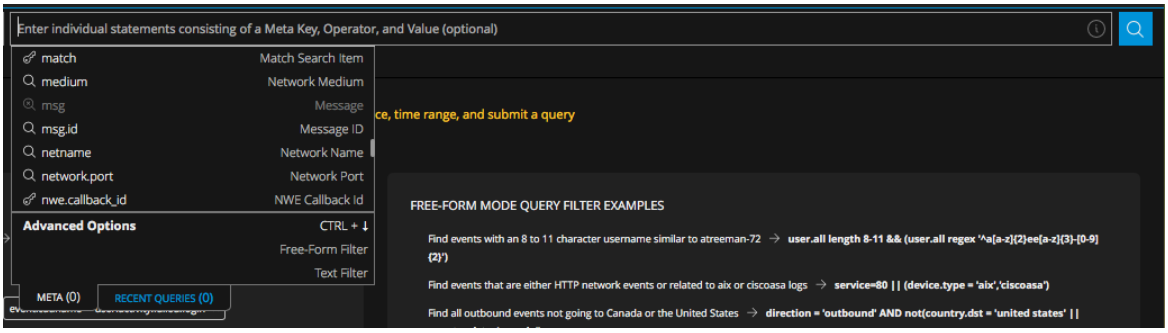
Add a Simple Filter in Guided Mode

To create a simple filter in Guided Mode:

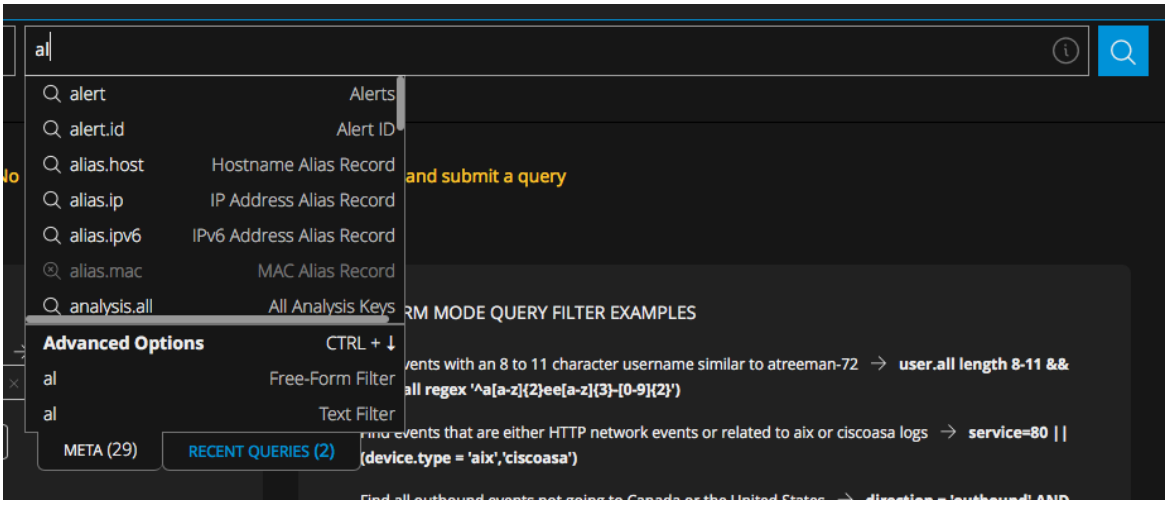
1. Go to the **Events** view (Event Analysis view in Version 11.3 and earlier) and do one of the following:
   a. (Version 11.4.1 and later) Click in the query bar and when the filter entry form is displayed, select the Meta tab if it is not already selected.
   b. (Version 11.4 and later) Select **Guided Mode**, click in the query bar and when the filter entry form is displayed, select the Meta tab if it is not already selected.
   c. (Version 11.2 and later) Select **Guided Mode** and click in the query bar.
   d. (Version 11.1) Click in the empty query bar, or before or after an existing filter. This is an example of the empty query bar in Guided Mode before you begin entering a filter.



If the insertion point is between two filters, a green circle (Version 11.3 and earlier) or a bold cursor (Version 11.4 and later) marks the insertion point. If the insertion point is at the end of the query bar, the filter entry form opens with a blinking cursor at the entry point. A drop-down list displays the available meta keys passed from the service being investigated in alphabetical order. This figure shows the filter entry form from Version 11.4.



2. To select a meta key do one of the following:
   a. If there is only one option in the drop-down list, press **Enter**.
   b. If there are two or more options in the drop-down list, click a meta key or select a meta key using the up/down arrows, then press **Enter**.
   c. Start typing the meta key. As you type the meta key, the list is filtered to include only meta keys that contain the text you typed. The count next to the label on the Meta (0) tab increments to enumerate the indexed meta keys that match the typed text. Keys that are not indexed are disabled and not selectable and are not included in the count, for example, `alias.mac` in the figure below is not indexed and is dimmed. Click a meta key or select a meta key using the up/down arrow, then press **Enter**.

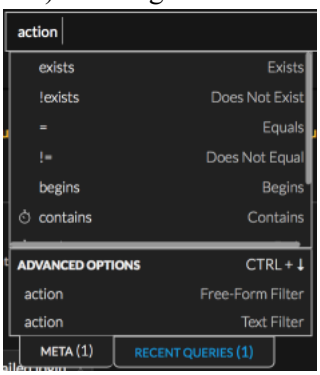


   d. To select a highlighted meta key, press **Enter**.
   The count on the Meta label changes to 1.

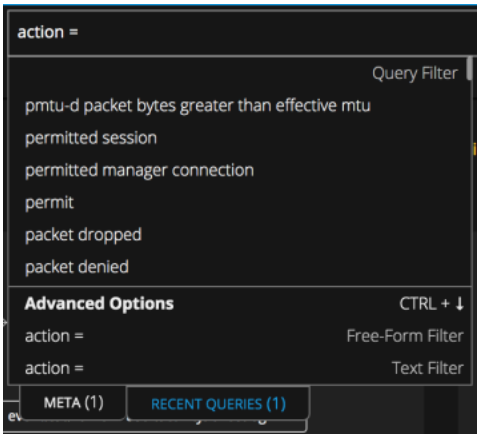
> **Note:** If no meta key in the drop-down list is selected, and the list has no meta keys to select, either the Free-Form Filter or the Text Filter option is highlighted based on the content already typed in the query bar.
> --If the text typed in the query bar includes some form of query syntax and other operators not yet supported by the user interface, the Free-Form Filter option is highlighted and you can create a free-form filter. In Version 11.3 and earlier, the `**`, `&&`, `||`, `()`, `AND`, `OR`, `comma`, `-`, `length`, and `regex` operators are not supported by the user interface. The Version 11.4 user interface supports these operators. If the Free-Form Filter is not highlighted, and the query bar has no text filter, the Text Filter is highlighted so you can create one.
> --If the first condition is true, and there is already one text filter, the Free-Form Filter option is highlighted so you can create a free-form filter.

   e. If you want to edit or delete the meta key, press **Backspace** or **Delete**.
   As you backspace and delete characters, the meta key drop-down list is filtered to include meta keys that contain those characters. To select a meta key, press **Enter**.
   The meta key is added to the filter entry form, and a list of valid operators for the selected meta key is displayed. Operations that require more time to process are marked by a (stopwatch

icon). This figure shows the stopwatch icon marking the `contains` operator.

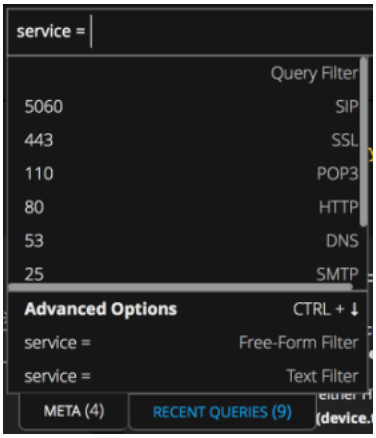


3. To select an operator, do one of the following:
   a. If there is only one option in the operators drop-down list, press **Enter** to select it.
   b. If there are two or more options in the operators drop-down list, click an operator or select one using the up/down arrows, then press **Enter**.
   c. Type the operator and press **Enter**. As you type, the operators drop-down list is filtered to show only operators that contain the typed text. Click an operator or select one using the up/down arrows, then press **Enter**.
   The operator is added to the filter entry form. In Version 11.4 and later, if the operator accepts a value, the suggested values drop-down list is displayed. Earlier versions leave the cursor in the filter entry form so that you can type a value.

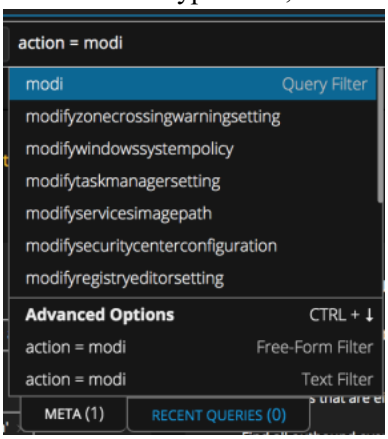


4. (Optional) If the selected operator in the filter entry form accepts a value, do one of the following:
   a. In Version 11.3 and earlier, type the value and press **Enter**.
   b. In Version 11.4 and later, paste a value that you have copied from somewhere and press **Enter**.
   c. In Version 11.4 and later, begin typing in the **Query Filter** field.
   As you type, the meta value drop-down list is filtered to return up to 100 properly indexed values that begin with the typed text. The suggested values are based solely on the time range; filters in the query do not filter the list of 100. The auto-suggest function looks for matches in all events in the current data set, not just the (up to 10,000) downloaded events. If nothing in the list matches exactly, the text you typed in the Query Filter field is highlighted and this message tells you that

no suggestions were found. Some values, such as the integers for the `service` meta key, also display the definition of the service type.



If there is an exact match, that value is highlighted. In the following example, there is no exact match for the typed text, modi.

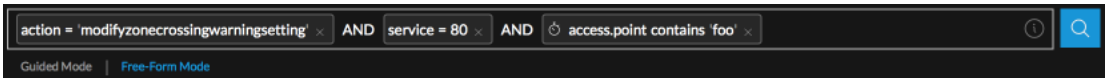


i. If the typed text is the value you want to use in the filter, press **Enter**.

ii. If you see the value that you want to query in the list and it is not highlighted, click the value or use the up/down arrows to highlight the value. Then press **Enter**.

iii. If you want to edit or delete the value, press **Backspace** or **Delete**.
As you backspace and delete a character, the meta value drop-down list is filtered to include values that begin with the remaining characters. To select a value, press **Enter**.
The value is added to the filter entry form.

5. To create the filter, press **Enter**. If you click anywhere outside the box before pressing **Enter**, the filter is not created.
The new filter is inserted, and the blinking cursor is refocused after the last filter, the meta keys drop-down list is displayed. If there is an error in the filter, it is outlined in red. You can hover over the filter to see a tool tip explaining the error. This figure shows a query being created with no errors.



6. If the filters have no errors, you are ready to execute the query in the query bar. Click  .
The results are returned and loaded in the Events panel. The first 10,000 events that match the query

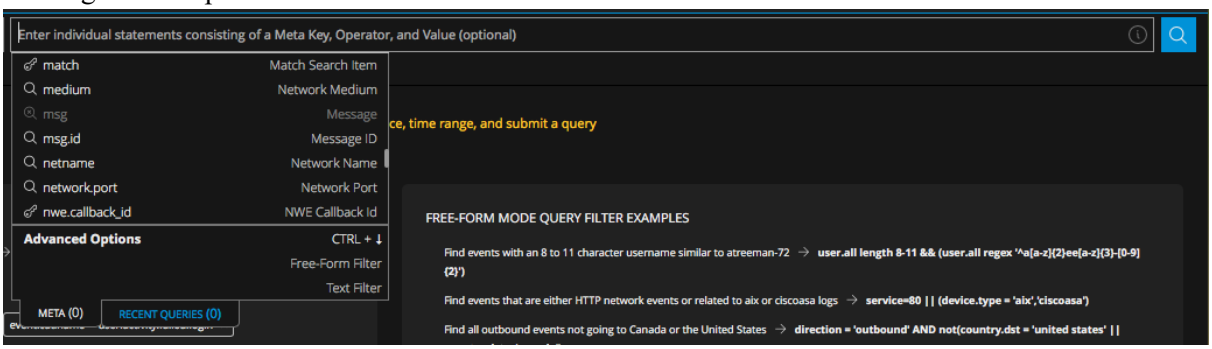
begin loading in the Events panel. As the events are loaded, a status bar at the top tracks progress and you can scroll to the bottom of the list to see the completion status.

7. (Optional in Version 11.3 and later) If you want to see detailed status in the Query Console, click the information icon .
8. (Optional in Version 11.3 and later) If you want to cancel the query before it finishes executing, click .
   The query stops executing and a notification that the query has been canceled is displayed.

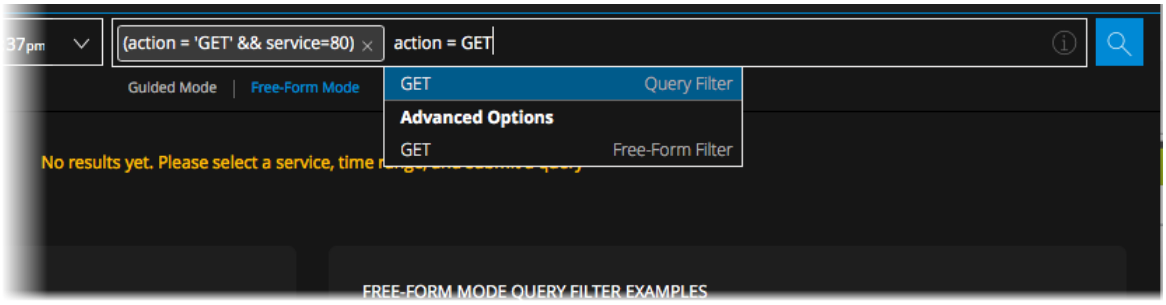
Add a Free-Form Filter in Guided Mode (Version 11.3 and Later)

To filter the data displayed in the Events view using a free-form filter in Guided Mode:

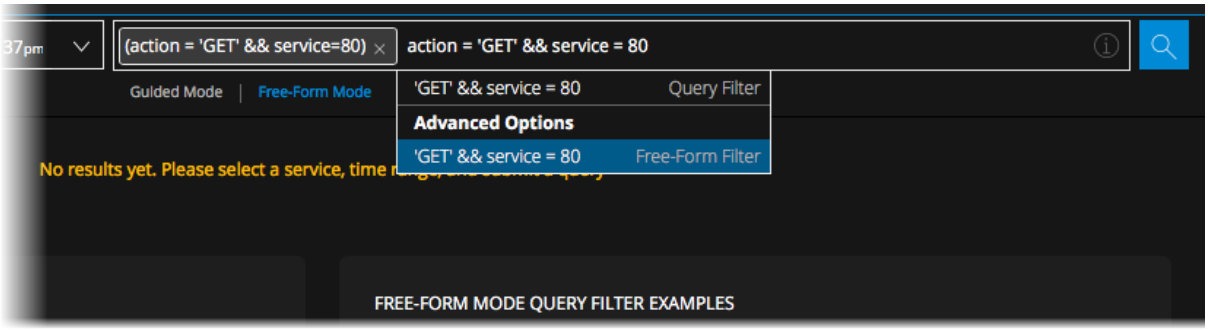
1. Go to the **Events** view, select **Guided Mode** below the query bar, and click in the query builder field. (For Version 11.4.1, simply click in the query builder field.)
   If the insertion point is between two filters, a green circle or a bold cursor marks the insertion point. If the insertion point is at the end of the query bar, the filter entry field opens with a blinking cursor at the entry point. A drop-down menu lists available meta keys passed from the service being investigated in alphabetical order.



2. Do one of the following:
   a. Place the cursor in the **Free-Form Filter** field and begin typing the query.
   b. Begin typing the filter beginning with a meta key or with an open parenthesis. When entering and editing filters in the query builder, parentheses pairs are automatically balanced. If you type an open parenthesis, the other part of the pair is added to the filter.
   When no matching meta keys or operators are available in the drop-down menu, the Free-Form Filter option becomes available, and the text you typed is available in the **Free-Form Filter** field.

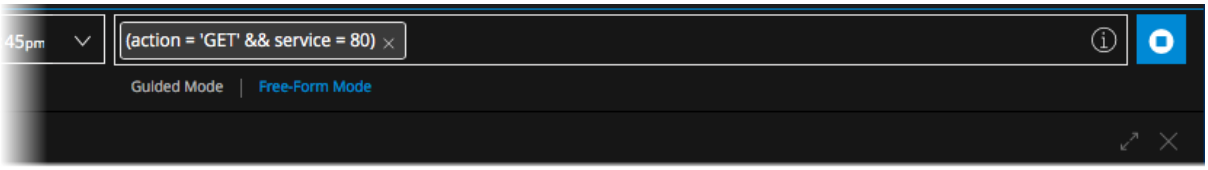


3. Continue typing the entire expression and press **Enter**.
   (If you click anywhere outside the box before pressing **Enter**, the filter is not created.) This figure shows a free-form expression created by continuing to type after the value GET.



   The new filter is inserted, and the blinking cursor is refocused after the last filter, a new filter entry form is displayed. If there is an error in the filter, it is outlined in red. You can hover over the filter to see a tool tip explaining the error.

4. To execute the query, click . While the query is executing, the button changes to .



5. If you want to cancel the query before it finishes executing, click .
   If you do not cancel the query, you can click to view the status of query execution. When the query is finished executing, the Events panel displays appropriate results for the query.

Add a Text Filter to Find a Value Anywhere in the Data Set (Version 11.4 and Later)

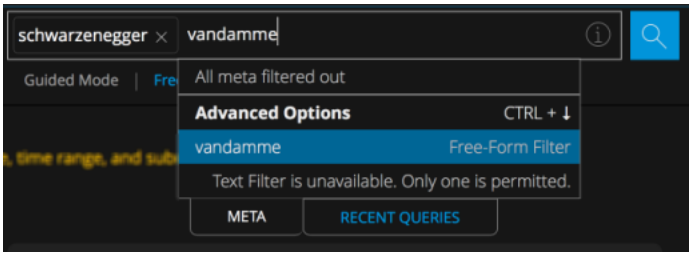
In Version 11.4 and later, the text filter allows you to find a specific value in the current data set (endpoint, logs, and network events). The text filter initiates a case-insensitive search against all the data for meta keys that are indexed by value. The text filter does not search for values that are indexed by meta key or not indexed so you not see all results. A message advises that `Results may be limited by a text filter, which matches only indexed meta keys. If you want to conduct a more exhaustive search against raw events, click here and choose the appropriate options in the Search Events drop-down menu.` Icons in the drop-down list indicate the index level of each meta key:

- filename.size - indexed by meta key
- filename.src - indexed by meta value
- float32.whatever - not indexed

Note: All services in the hierarchy being queried (Broker, Concentrators, and Decoders) must be at Version 11.3 or later. The text filter is not available in the drop-down menu when there are services below Version 11.3 in the hierarchy.

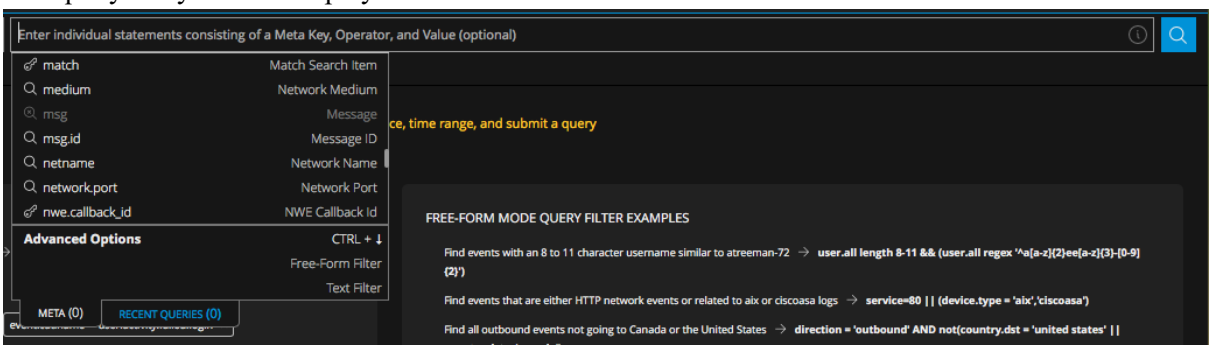
The text filter is useful when you have some idea of what you are looking for, but are not sure where to look (which meta key or service). As an example, if you are interested in looking for a file name, click in the query bar, type the complete text string, and click **Text Filter**. The text filter initiates a search against all the data in the index, within the services and time range being investigated, and returns exact matches to the text string.

A query can include one text filter and any combination of simple and free-form filters. The operator for a text filter must be AND because it acts as a filter over the results of all the other filters in the entire query. If one text filter already exists in the query bar, the Text Filter option is disabled as shown in the figure below. Text filters are not supported inside parentheses.

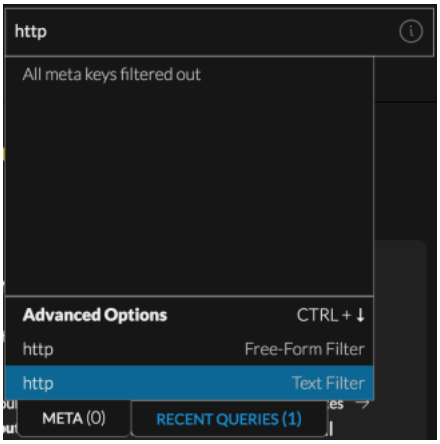


To create a text filter:

1. Go to the **Events** view and click in the query bar.
   The query entry form is displayed.

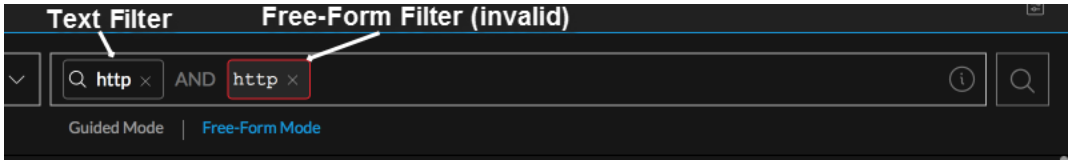


2. Type the text string that you want to find, for example, `http`.
   The text string is displayed in the meta key drop-down list under Advanced Options.

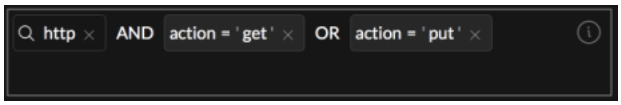


3. Click **Text Filter** under **Advanced Options**.
   The text filter is created in the query bar. The following figure illustrates the different appearance of

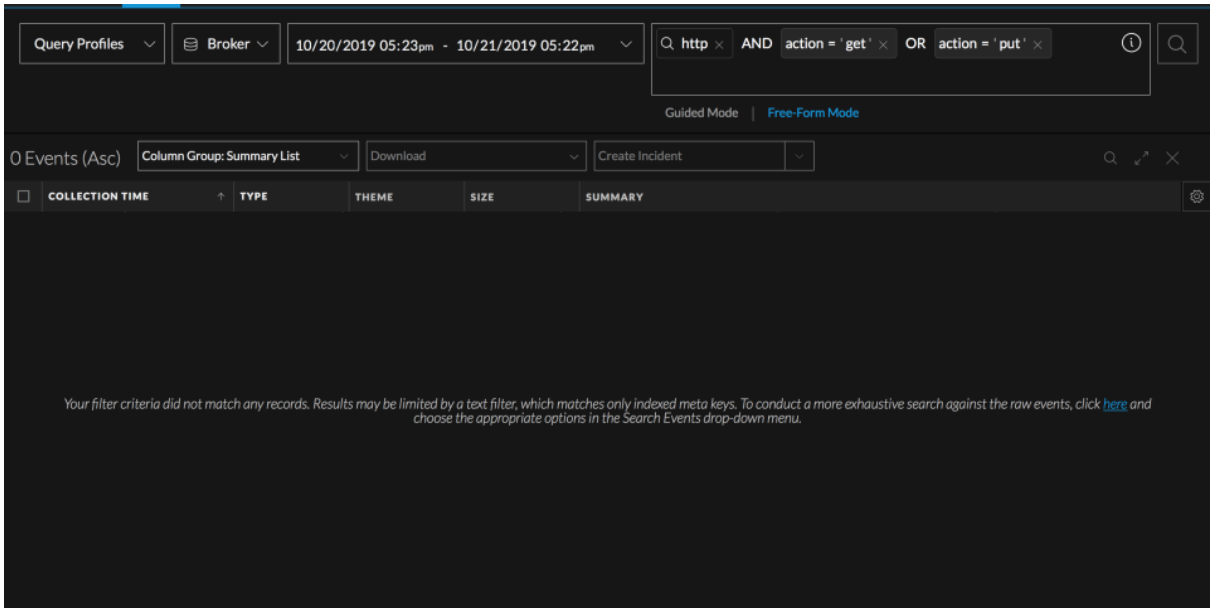
a text search filter versus a free-form filter. The free-form filter is in a fixed-space font and outlined in red. The red outline indicates a syntax error because a valid expression is expected in a free-form filter. The text filter is marked by the search icon. No syntax requirements are applied to text search filters.



4. (Optional) Create additional simple or free-form filters in the query bar. There can be only one text filter in the query. This example was created by typing `http` as a text filter and then adding two more filters - `action = 'get' OR action = 'put'`

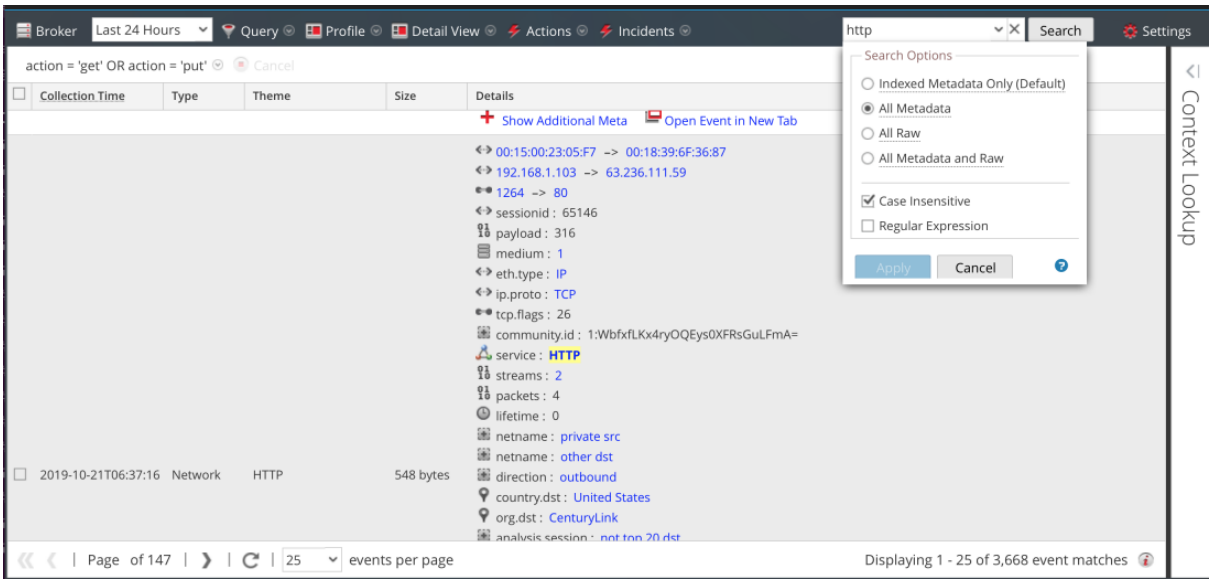


5. To submit the query, click  .
The results are displayed in the Events panel. This figure illustrates the Events panel with no results displayed and a message with instructions for improving results. Every time you use a text filter, this message is at the bottom of the results offering a link to expand your search.

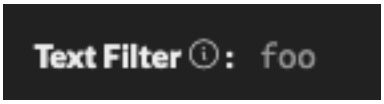


6. Click the **here** link in the message.
A new browser tab opens with the query results displayed in the Legacy Events view, where you have additional options to improve the search. This figure shows the results for the same query when

non-indexed metadata is included.



7. Click the Information icon in the query console to view the status of the query. This figure shows a text filter in the query console.



Select All Filters and Copy All Filters in the Query Bar (Version 11.4.1 and Later)

While creating a filter in the Events view query bar, you can use keyboard commands to select all filters (Ctrl-A for Windows OS, Cmd-A for MacOS) and then copy the selection to the local clipboard (Ctrl-C for Windows OS, Cmd-C for MacOS).

To select all filters and copy them to the clipboard:

1. In the **Events** view > **Events** panel, click on a focused pill or in the query entry form, and press **Ctrl-A** for Windows OS or **Cmd-A** for MacOS.
   All filters in the query bar are selected.
2. To copy the selected filters to the clipboard, type **Ctrl-C** for Windows OS or **Cmd-C** for MacOS.
   You can share the clipboard with other analysts or paste the contents in the query bar.

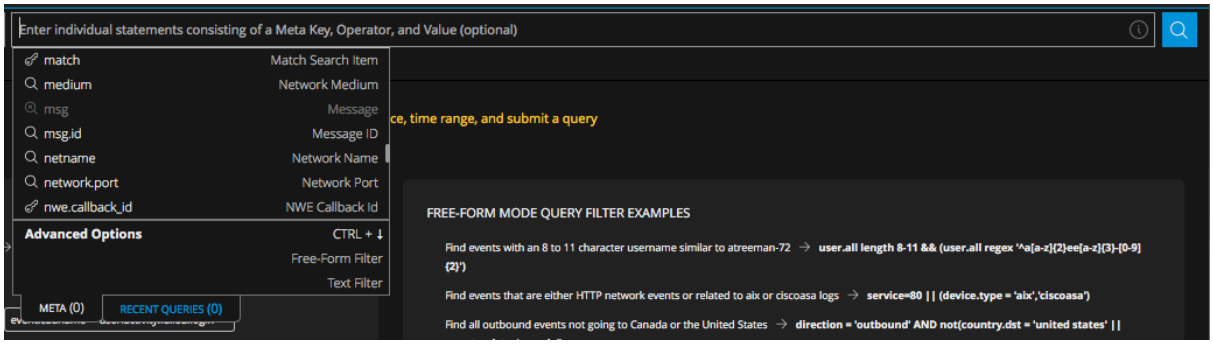
Paste Text in the Query Bar (Version 11.4 and Later)

While creating a filter in the Events view query bar, you can paste instead of typing the complete text of a filter that you have copied from somewhere else. You can paste the text into an empty query bar or next to an existing filter in the query bar. Depending on the text you typed, the query parsing engine parses the information that you pasted and creates a new filter, which can be a simple filter, a free-form filter, or a text filter.

- A text string of this form is added as a new simple filter in the query bar: `<valid meta key> <valid operator> <optional value>`. This is an example: `alias.host contains 's'`.
- A text string of this form is added as two simple filters in the query bar: `<valid meta key> <valid operator> <optional value> && <valid meta key> <valid operator> <optional value>`. This is an example: `alias.host contains 's' && action exists`, which is converted to `alias.host contains 's' AND action exists`.
- A text string that contains unparsable text may be converted to a free-form filter. For example, using `NOT (device.ip = 10.10.10.10)` is unsupported for creation of a filter in Guided Mode, so this would be converted to a free-form filter. Free-form filters are validated by the server when they are submitted.
- Text that does not conform to the filter syntax is added as a free-form filter.

To create a filter by pasting text:

1. Go to the **Events** view > **Events** panel, select **Guided Mode** under the query bar, and click in the query bar. (For Version 11.4.1, simply click in the query bar.)
   The query entry form is displayed.



2. Type **Ctrl-V** (Windows OS), **Cmd-V** (MacOS), or **right-click** and **Paste** to paste text that you have copied into the clipboard from somewhere else. Do one of the following:

   a. If the text you pasted is a statement that can be parsed, one or more simple filters is created.
   If the text you pasted is a statement that cannot be parsed, a new free-form filter is created.
   If the text you pasted is not a statement and not a valid meta key, an invalid syntax error is displayed.
   If you pasted a valid meta key for a new filter you are building, the meta key is highlighted in the drop-down list, and you can continue creating a filter as usual by entering an operator and a value.
   After you select a valid meta key and a valid operator (for example, `city.dst =`) any text that you paste is treated as a text string if the meta key supports a text value, and one filter is created.
   If the meta key does not support a text value all of the text in the query bar is parsed as described in step a above.

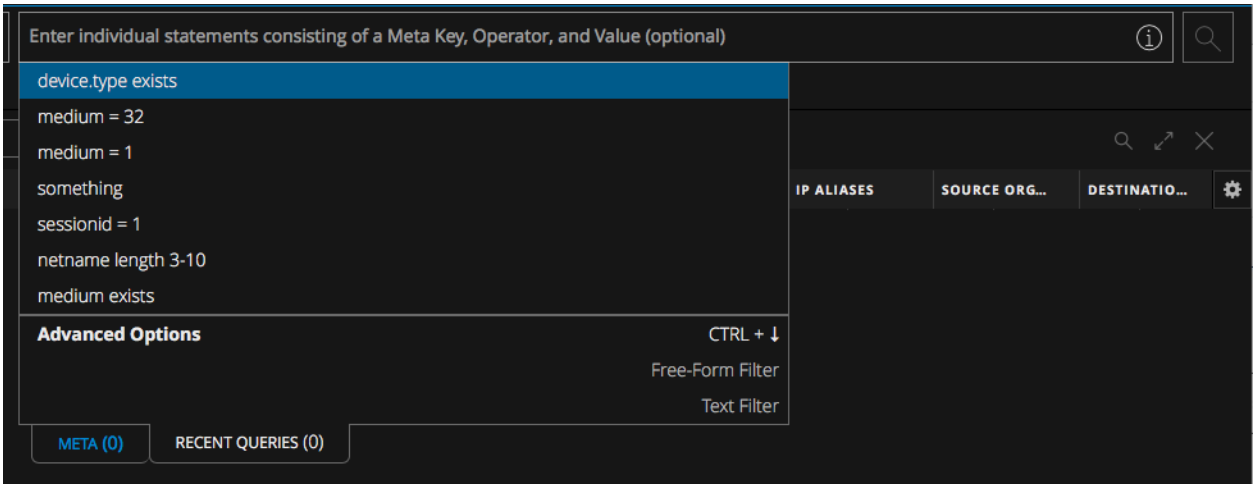
3. Add more filters in the query bar if you wish, and then submit the query.
   The query is executed.

Insert a Filter Based on a Recent Query (Version 11.4 and Later)

In the Guided Mode query bar, you can insert a filter based on a recent query. When the Recent Queries tab is opened and nothing has been typed in the query bar, up to 100 of your most recently executed queries are displayed in a scrollable list. The list is sorted to show the most recent at the top, and the Recent Query count is set to 0. When you begin typing, the list is filtered to display up to 100 queries from the query history database that contain matching text, even if the matches are not in the most recent 100 queries. The Recent Query count changes to reflect the number of matching queries as you type.

The top entry in the list is highlighted by default. To select a recent query, you can move the highlighting up and down in the list using the up and down arrow or by mouse-over of a recent query. As you type the list is filtered and the highlighting moves back to the top of the list. Clicking a query, or pressing Enter while a query is highlighted, creates a new filter with the text of the selected query.

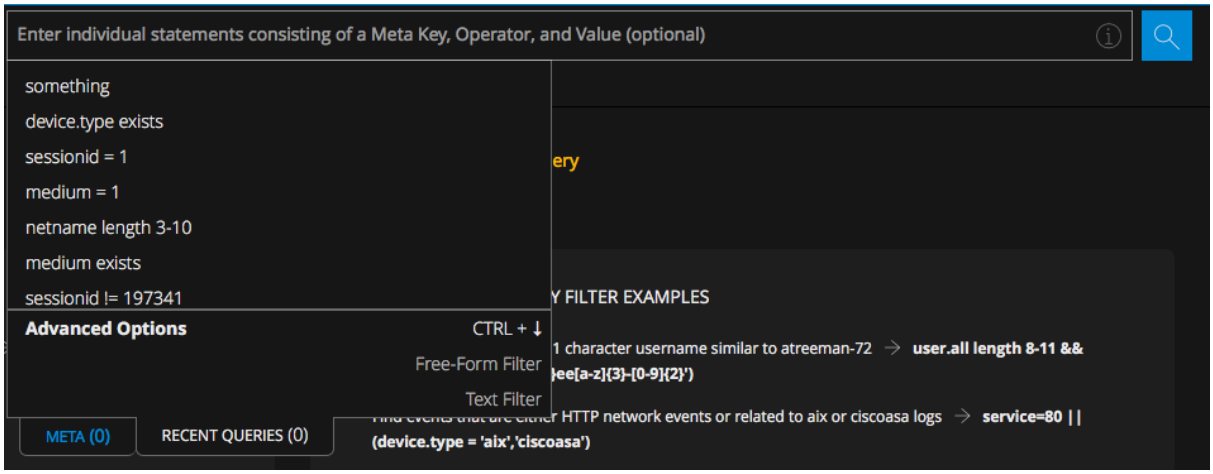
Whenever you submit a query, the list is sorted to add that query, now the most recent, at the top.



To create a filter based on a recent query:

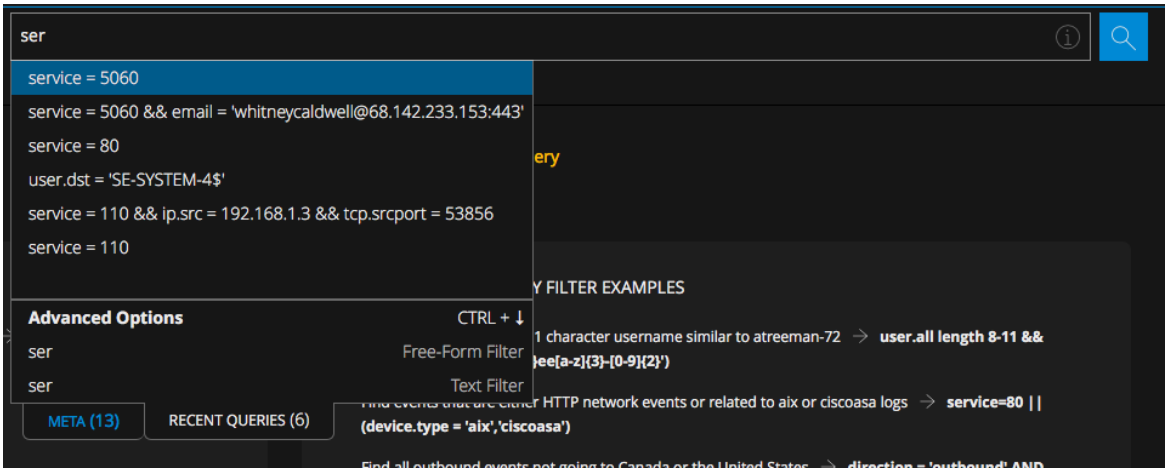
1. Go to the **Events** view, select **Guided Mode** under the query bar, and click in the query bar. (For Version 11.4.1, simply click in the query bar.)
   The Meta Key drop-down list is displayed in the Meta tab.

2. Select the **Recent Queries** tab.
   The Recent Queries drop-down list is displayed with a count of 0.



3. To search for a recent query, do one of the following:

   a. Begin typing some text.
      As you type more characters or backspace to delete characters, the list is filtered to show recent queries that contain the text you typed. The count in the Recent Queries label increments to show the number of matching queries as you type.



   b. To select a query and add a new filter, continue to type and use the up and down arrows until the query you want to use as a new filter is highlighted.

   c. With a query highlighted, press **Enter** or simply click a query that you see in the list.
      The filter is added in the query bar.

4. Add more filters in the query bar if you wish, and then submit the query.
   The query is executed and the list is sorted to add that query, now the most recent, at the top.

Edit a Filter in Guided Mode

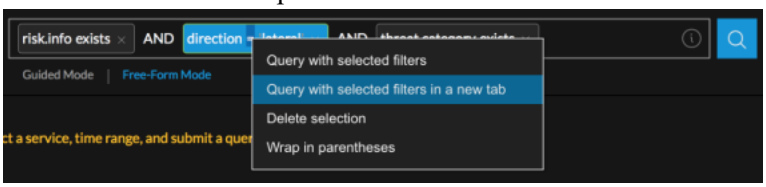
With a query in the Guided Mode query bar, you can edit a filter. To edit a filter:

1. Double-click the filter, or click the filter and press **Enter**.
2. Edit the filter. When finished editing, press **Enter** to update the filter.
3. If you want to execute the query again, click .
   The Events panel displays results for the updated filter.

Query Using Selected Filters in Guided Mode

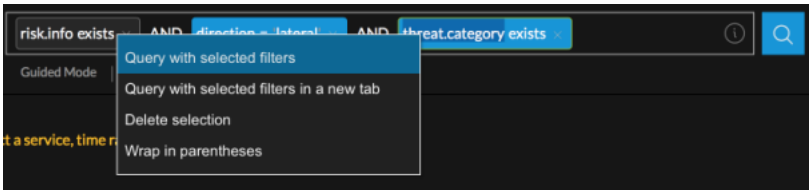
When you have one or more filters in the query bar in Guided Mode, you can refocus the query to include only selected filters, displaying results in the current browser tab or a new browser tab. Some filters include expressions with nested parentheses in Version 11.4, and you can refocus part of a filter that includes nested parentheses. To update the query using only selected filters, do one of the following:

1. Using a query that includes one or more simple filters, for example a query has three filters: `risk.info exists`, `direction ='lateral'`, and `threat.category exists`.
   a. Select `direction = 'lateral'`, right-click the filter and select **Query with selected filters in a new tab** in the drop-down menu.



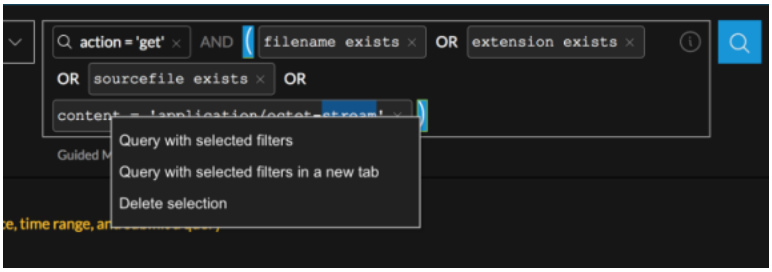
   A new tab opens with the results for the selected filter and the original query is left intact on the previous tab.

   b. To query the selected filters in the same tab, select `direction = "lateral"` and `threat.category exists`. Then right-click and select **Query with selected filters** in the drop-down menu.



   A query with only the selected filters is submitted and all remaining filters are removed.

2. (Version 11.4) For a query that includes a filter containing nested parentheses, for example: `action = 'get' AND (filename exists OR sourcefile exists OR content = 'application/octet-stream')`, do one of the following:



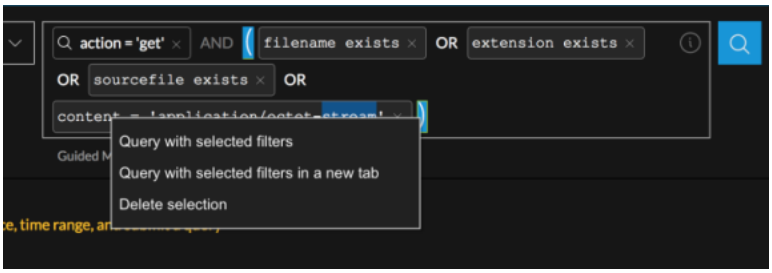
a. Select the close parenthesis after `'application/octet-stream'`, right-click, and select **Query with selected filters in a new tab**.
A new tab opens with results for `(filename exists OR sourcefile exists OR content = 'application/octet-stream')`.

b. Select the same, right-click, and select **Query with selected filters**.
The results for `(filename exists OR sourcefile exists OR content = 'application/octet-stream')` are displayed in the current tab.

Delete a Filter and Delete Text or Parentheses in a Filter in Guided Mode

Some keystroke editing features became available in Version 11.4; these are labeled in the steps.

1. To delete a filter, do any of the following:
   a. Click **X** in a filter.
   b. Select the filter and press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS).
   c. (Version 11.4 and later) Select the filter and press **Backspace** (Windows OS) or **Delete** (MacOS).
   d. Right-click one or more filters and select **Delete selected filters** or **Delete selection** (Version 11.4 and later) in the drop-down menu.
   The filter and the operator to the right or left of the filter is deleted, ensuring that no extraneous operators remain in the query bar.
2. (Version 11.4 and later) To delete characters in a filter or parentheses and contents in a filter, do any of the following:
   a. To delete the previous character: With the cursor next to a character in the query bar, press **Backspace** (Windows OS) or **Delete** (MacOS).
   b. To delete all characters: With the cursor in a filter, press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS).
   c. To delete the selected characters: With characters selected in the query bar, press **Delete** or **Backspace**.
   d. To delete parentheses, but not the characters inside the parentheses, select one of the parentheses and press **Delete** (Windows OS) or **Fn** + **Delete** (MacOS).
   e. To delete a set of parentheses and the contents, for example, `(filename exists OR sourcefile exists OR content = 'application/octet-stream')`, select the parenthesis after `get`, right-click, and select **Delete selection**.

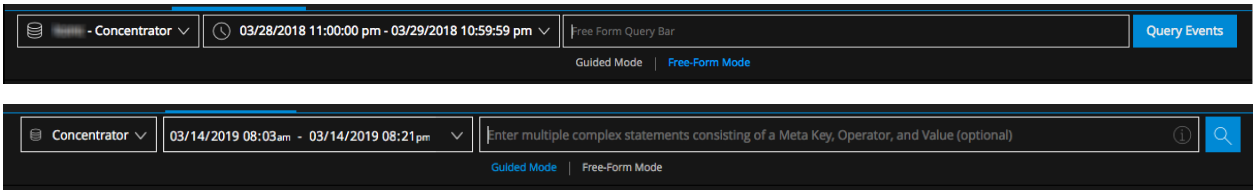


Everything except `action = 'get'` is deleted.

Create a Query in the Free-Form Mode

Free-Form Mode is used in Version 11.2, 11.3, and 11.4, but no longer available in Version 11.4.1.

Free-form queries are most useful when you have a long text string saved that you want to paste, or if you have one in mind that you want to enter quickly, and you know the meta keys, valid operators, and valid syntax for entering values. The following figure illustrates the initial Events view with the empty Free-Form query builder field. The first example is Version 11.2 and the second example is Version 11.3.



The blinking cursor indicates that you can enter a query. You can enter free text here. As more expressions are added and they cannot be displayed in a single line, they wrap to another line and the input area expands vertically so that all filters are visible without scrolling to the right.

These are some examples of queries that you can enter in Free-Form mode:

To find events with an 8- to 11- character username similar to atreeman-72:

```
user.all length 8-11 && (user.all regex '^a[a-z]{2}ee[a-z]{3}-[0-9]{2}')
```

To find events that are either HTTP network events or related to aix or ciscoasa logs:

```
service=80 || (device.type = 'aix','ciscoasa')
```

To find all outbound events not going to Canada or the United States:

```
direction = 'outbound' AND not(country.dst = 'united states' || country.dst = 'canada')
```

If you have a submitted query in Guided Mode, the query is transformed into text when you click switch to Free-Form mode. This is an example of a query submitted in Guided Mode as two filters, `service = 80` and `direction = 'outbound'`, and then viewed in Free-Form mode.



The button on the right side of the query builder is visible as needed to input a query. The query is applied when you click . At that time the query is validated to show syntax and logic errors.

Operations that require more processing time are not highlighted as they are in Guided Mode, but this table provides a summary of expensive operations for reference.

Index Method	Non-Text Value	Text Value	Regular Operations	Expensive Operations
By Key	✓		`exists`, `!exists`	`eq`, `!eq`
By Key		✓	`exists`, `!exists`	`eq`, `!eq`, `begins`, `ends`, `contains`
By Value	✓		`exists`, `!exists`, `eq`, `!eq`	no expensive operators
By Value		✓	`exists`, `!exists`, `eq`, `!eq`, `begins`	`ends`, `contains`
By None	special case for `sessionid`		`exist`, `!exits`, `eq`, `!eq`	no expensive operators

Filter Results in the Navigate View

When conducting an investigation in the Navigate view, there are several methods available to refine the results displayed when meta key values are loaded in the Navigate view. The rest of this topic is focused on the basic methods of filtering data:

Note: By default, the Navigate view is disabled in Version 11.6 as the Filter Events Panel in the Events view provides this functionality. To enable the Navigate view, see Configure the Navigate View and Legacy Events View.

- Set the Time Range
- Set the Quantification Method and Sort Sequence of Meta Key Results
- Manage and Apply Default Meta Keys in an Investigation
- Drill into Data in the Navigate View Time Chart
- Drill into Data in the Values Panel

Set the Time Range

When conducting an investigation in the Navigate view, the time range options limit the results returned. You can select:

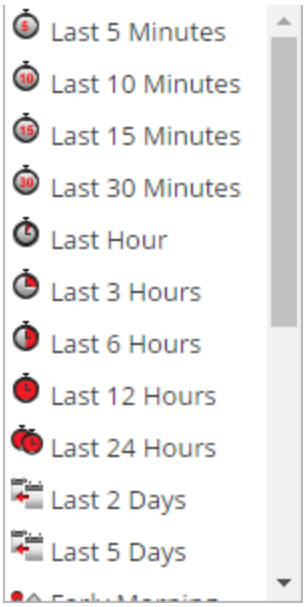
- A time range relative to the collection- ranges relative to the collection are based on the last collection time for data.
- A time range relative to the calendar.
- A custom date range.
- All data.

The selected Date Range is shown in the Navigate view tool bar as the Time Range label. By default the label is **Last 3 Hours**. The Time Range displayed in the timeline banner shows the first and last timestamp for the date range being used for the metadata.

Note: Time range is based on the Time Zone configured in the Profile Preferences panel as described in "Setting User Preferences" in the *RSA NetWitnessGetting Started Guide*.

To select a built-in time range:

1. Click the **Time Range** option in the Navigate view toolbar. The default time range is for the **Last 3 Hours**, but a different value from the selection list, for example, **All Data** or **Last Hour,** may already be selected and used as the label in the options panel.
   The Time Range selection list is displayed.



2. Do one of the following:
   - If you want to see all data, select **All Data**.
   - If you want to set a time range in minutes, hours, or days that is relative to the collection, select a value such as **Last 10 minutes**, **Last 3 Hours**, or **Last 5 days**.
   - If you want to set a time range relative to today, select **Yesterday**, **This Week**(Version 11.1), **Last Week** (Version 11.1), **All Day**, or a part of the day such as **Early Morning**, **Morning**, **Afternoon**, or **Evening**.
   - If you want to set a unique date range, select **Custom** in the **Time Range** menu and follow the procedure below.
     The selected time range is applied to the current results in the Values panel.

To specify a custom time range:

1. Select **Custom** in the **Time Range** menu.
   Date selection options are displayed in the toolbar.



2. Within the time **Start Date** and **End Date** fields, do the following to specify the date and time:
   a. Click a date from the calendar.
   b. (Optional) Select the time from the Hour and minute fields or click **Now**. The time selection defaults to the current time of day.

Note: The value for start time in seconds always defaults to :00, and the value for end time in seconds always defaults to :59. For example, if you are using time to drill down into an issue, the drill time is interpreted as "HH:MM:00 - HH:MM:59."

3. To apply the range, click **Go**.
   The selected time range is applied to the current results in the Values panel.

Set the Quantification Method and Sort Sequence of Meta Key Results

You can select the way results for each meta key are quantified and sequenced in the Navigate view.

Note: If meta entities (Version 11.1 and later) are used in meta groups, the results will show the top 20 values that matched any of the meta keys contained in the meta entity.

Each meta key section in the Navigate view contains an ordered list of values showing each meta key value (Value) and its count (Total). You can specify whether:

- The results in each meta key section are sorted based on Value or Total.
- The results are sorted in ascending or descending order.
- The values shown for each meta key are quantified by number of packets (Packet Count), number of sessions or logs (Quantify by Event Count) or by the size of events (Quantify by Event Size).

Note: If you have both a log decoder and a packet decoder for which you are viewing the metadata, the calculation of what is actually being counted is dependent on the type of key. If you select to Quantify by Packet Count and are looking at logs, the Navigate view output is the same output as if you had selected Quantify by Event Count (see Navigate View for details).

This image shows the `Event Type` meta key presented in order by **Total** in **Descending** order. The value with the greatest count of matches is presented first. The value `failure audit` has 71 matches and is listed first. The value `logon` has only one match and is presented last. The quantification method is **Event Count**.



This image shows the `Event Type` meta keys presented in order by **Value** in **Descending** order. The value names are presented in alphabetical order starting at the end of the alphabet. The value `success audit` is listed first. The value `connect` is presented last. The quantification method is **Event Count**.

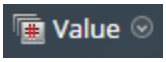


To select the quantification method of meta key count and ordering of meta key results displayed in the Navigate view:

1. In the toolbar, select **Event Count**, **Event Size**, or **Packet Count** and choose one of the quantification options in the drop-down menu. The label for the menu displays the selected option.
   

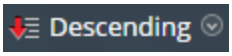
   The current view is reloaded according to your selection.

2. In the toolbar, select **Total** or **Value** and choose one of the ordering methods in the drop-down menu. The label for the menu displays the selected option.



   The current view is reloaded according to your selection.

3. In the toolbar, select **Ascending** or **Descending** and choose one of the sort order options in the drop-down menu. The label for the menu displays the selected option.
   The current view is reloaded according to your selection.



Manage and Apply Default Meta Keys in an Investigation

When analysts are conducting an investigation of captured data in Investigate, a default set of meta keys is loaded and displayed in a default sequence in the Navigate view > Values panel. The default content and sequence is based on the meta keys for the service being investigated. Analysts can specify the meta keys to display during navigation by selecting the default meta keys or by selecting a user-defined group of meta keys, which provides great flexibility to define meta keys. This can help to drill down more directly to the desired data and to reduce the load time by preventing the loading of meta that is not of interest in the current investigation.

Note: In Version 11.1 and later, wherever meta keys are used, you can also use configured meta entities.

If no custom meta groups are in effect, the Navigate view is displayed with the meta key visibility specified in the Default Meta Keys dialog. To optimize loading of meta keys in the Navigate view > Values panel, NetWitness does not open non-indexed meta keys by default. When you open a non-indexed meta key in the Values view, NetWitness begins loading values for that meta key. If the load time is excessive, the load of the meta key times out with a message. Title, values, and counts for non-indexed meta keys are not drillable in the Values panel. Additional labeling in Investigation identifies the non-indexed meta keys.

To select the meta keys to apply to your investigation, you can:

- Select the default meta keys.
- Select a set of meta keys, called a meta group.

Note: Investigate has built-in meta groups and user-defined meta groups. Once created, user-defined meta groups can be edited, deleted, exported for use on other services, and imported to the service you are investigating. All of these procedures are provided in a separate topic: Use Meta Groups to Focus on Relevant Meta Keys.

The Default Meta Keys dialog allows you to specify the default view and display sequence for meta keys during navigation in the Investigate > Navigate view for a specific service. For each key or for all keys, you can set the default view to:

- Hidden: Results for default meta key are hidden and are not available to load.
- Open: Results for default meta key are open with all values and counts displayed.
- Close: Results for default meta key are closed with only the meta name visible.

- Auto: The loading of default meta keys is controlled by the index level, which must be Indexed By Value.

When using the default meta keys, be aware that these can be modified for different services, and you may not be seeing the same set of default meta keys when navigating to a drill point on different services. If you do not see the expected data, you may need to change the initial view of the default meta keys.

When you change the initial state of default meta keys from within the Navigate view, the change persists for that service. When new keys are added to the custom index file for a Core service (for example, `concentrator-custom-index.xml` or `decoder-custom-index.xml`), the new keys are added to the default meta keys list. The changes made in the Navigate view apply only to the current service.

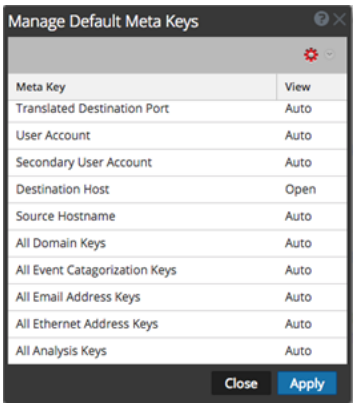
To specify that the initial Navigate view opens using default meta keys:

1. Go to **Investigate** > **Navigate**.
2. Select a service and select **Navigate**.
3. In the **Meta** menu, select **Use Default Meta Keys**.
   If an investigation is already in progress, the data is reloaded in the current view and an icon highlights the selected option. If no data is loaded yet, the default meta keys are used for the next load.

Configure Default Meta Keys

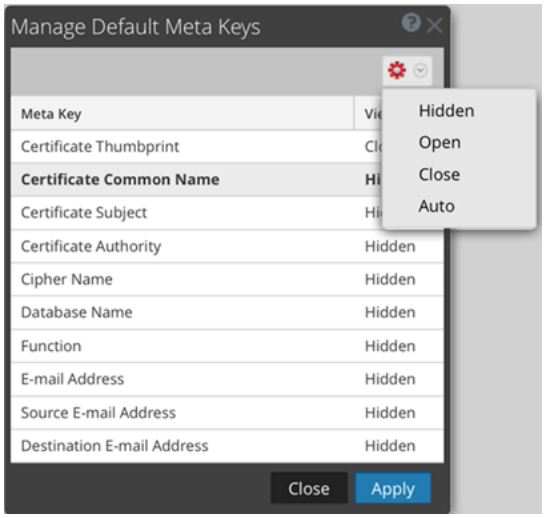
To configure the default view of default meta keys in the Navigate view:

1. In the **Navigate** view toolbar, select **Meta** > **Manage Default Meta Keys**.
   The Manage Default Meta Keys dialog is displayed with the list of available meta keys for the service.



2. (Optional) To change the order of the keys, select one or more keys, and drag the values up or down through the list of keys.
3. Do one of the following:
   - (Optional) To change the default view for all meta keys, make sure that no keys are selected and in the toolbar, select  .

- (Optional) To change the default view for one or more keys, select the keys and in the toolbar, select .
  A drop-down menu of possible initial views for all default meta keys is displayed.



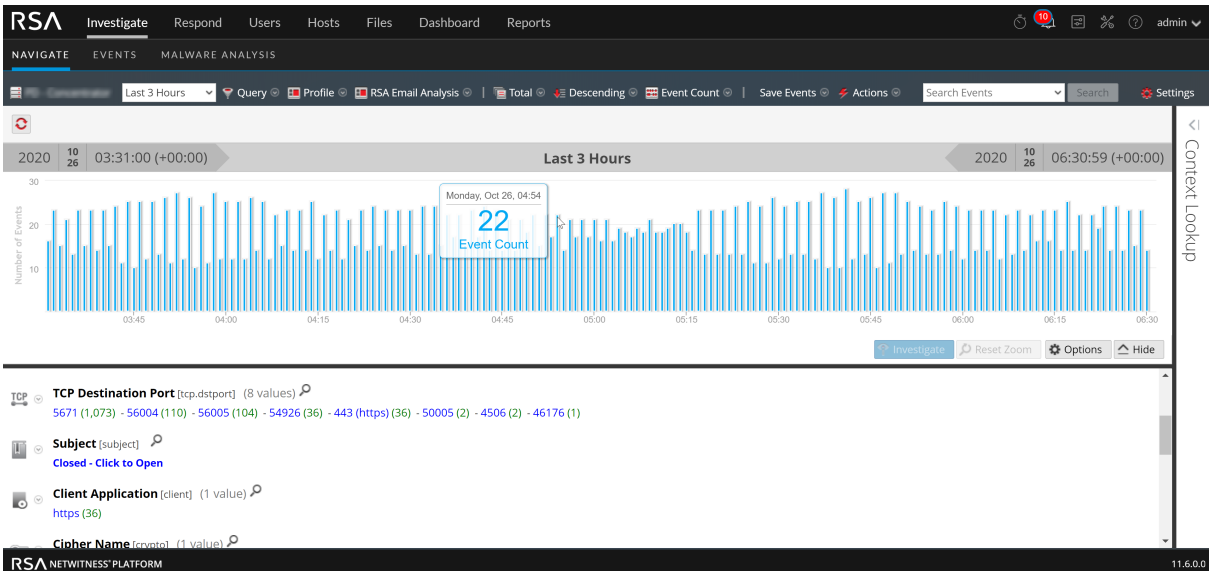
- (Optional) To revert to the default view for meta keys as specified in the service index file, make sure that no keys are selected and in the toolbar, select > **Auto**.
  When you modify the default view for a non-indexed meta key, you cannot set the key to OPEN. If you change the default view for a group of meta keys to OPEN and some of the meta keys are non-indexed, the non-indexed meta keys revert to AUTO. As a result, the meta key is automatically loaded only if it is indexed, and non-indexed meta keys are CLOSED until opened manually.

4. Select one of the views.
5. To save the changes, click **Apply**.
   The meta keys displayed in the Navigate view are set to your specifications. If the default meta keys are hidden, values for the meta keys are not shown in the investigation at all. If the default meta keys are closed, the values for the meta keys are not loaded by default, but you can load individual meta keys manually in the Navigate view.

Drill into Data in the Navigate View Time Chart

The Time Chart visualization allows analysts to visualize activity over time. You can zoom into the data by selecting a time window then selecting the Investigate option. You can then reset the navigation to the time range that was in effect before zooming.

1. Go to **Investigate > Navigate**.
   The Time Chart for the current drill point and selected time range is displayed. You can hover over the time chart to display total number of events occurred at a specific time.



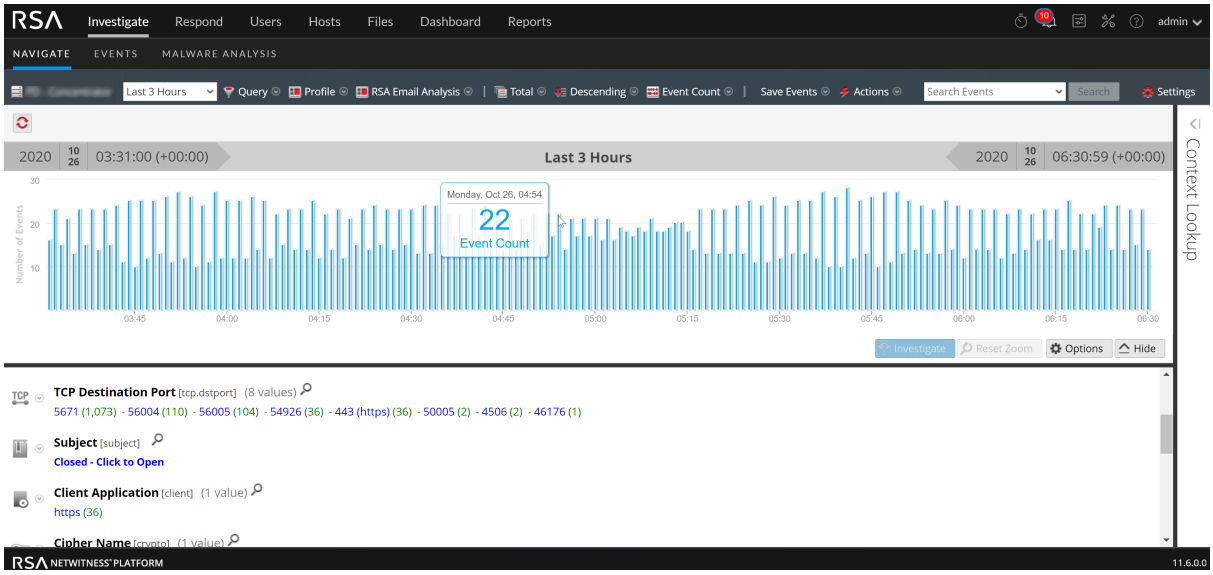
2. To highlight a period of time on the Time Chart, click over the desired time period and drag the mouse.
   The Time Chart is redrawn for the selected time range; however, the meta values are unchanged.
3. To drill into the data for the selected time range, click **Investigate**.
   The URL is updated to reflect the time range override, and the Investigation options panel is updated to reflect the custom time range. The Time Chart is redrawn and the meta values are loaded for the selected time range.
4. To reset the Time Chart to the original time range, click **Reset Zoom**.
   The URL is updated to reflect the original URL prior to zooming into the data, and the Investigation options panel is updated to reflect the time range selected before zoom. The Time Chart is redrawn for the selected time range and the meta values are loaded for that time range.

Drill into Data in the Values Panel

NetWitness displays the activity and values for the selected service in the Investigation > Navigate view. To investigate data, analysts drill into data by clicking on a meta key or a meta value, which is treaty as a query. In the Values panel, each query is added to the breadcrumb data in the Values panel. This results in a breadcrumb at the top with a crumb for each query. You can edit the breadcrumb to insert or remove a query.

To drill into a subset of the metadata:

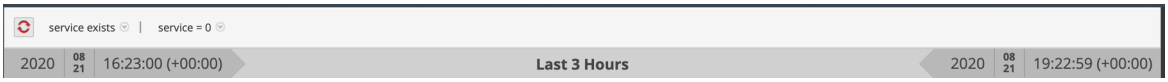
1. Begin an investigation so that metadata is displayed in the Navigate view.



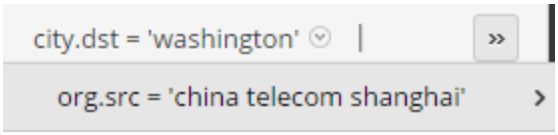
2. To drill down into the metadata, do any combination of the following:

   a. Click a **meta key**, for example, **Service Type**.

   b. Click a **meta value**, the blue text in the results. For example, **OTHER**.
   Each time you click a meta key or meta value, the investigation query pivots to a narrowed focal point, or drill point, in the data. At each drill point, the Values panel is updated and the new drill point is displayed in the breadcrumb. Below is an example of the first breadcrumb.



   This is an example of a long breadcrumb that does not fit in the toolbar. The last query that fits is followed by a drop-down menu that lists additional queries. To select a drill point within the overflow, click the overflow icon and a query in the drop-down list.

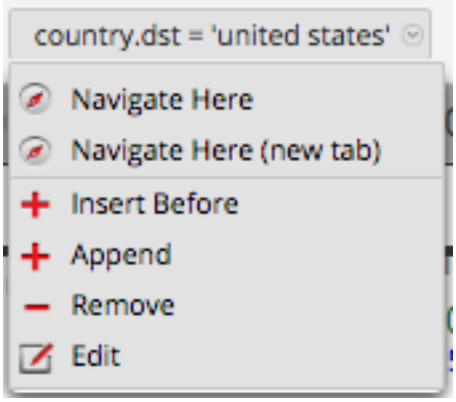


To add a query in the breadcrumb:

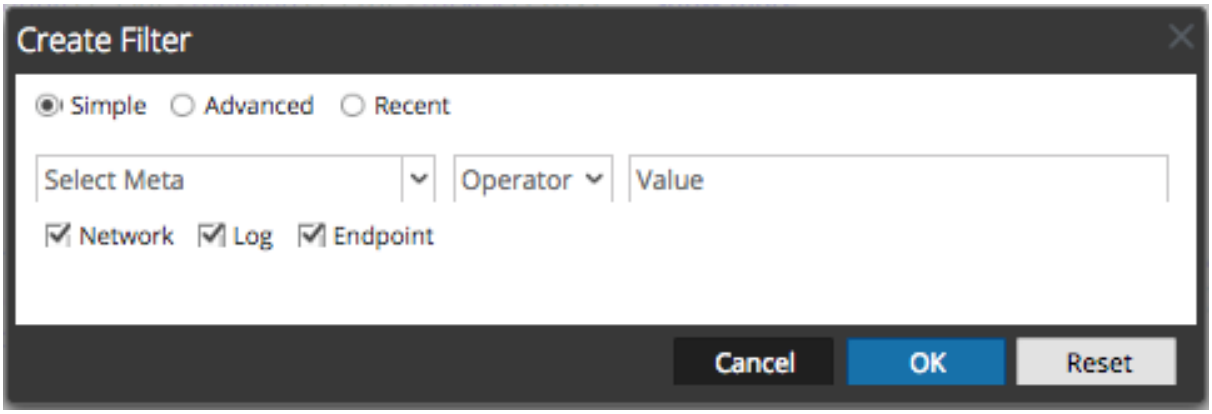
In the breadcrumb, you can click any of the crumbs to display the Query menu. You can insert a new query before a crumb, and append a new query to the end of breadcrumb. After each edit in the breadcrumb, NetWitness refreshes the results.

To add a query in the breadcrumb:

1. Click a crumb.
   The Breadcrumb menu is displayed.



2. To add a query in the breadcrumb, select **Append** or **Insert Before**.
   The Create Filter dialog is displayed.



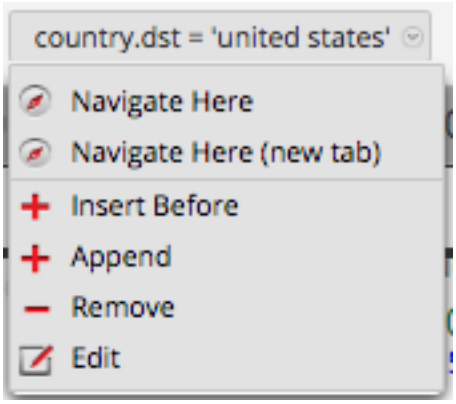
3. Create the Query as described in Create a Query in the Navigate and Legacy Events Views.

To edit a query in the breadcrumb:

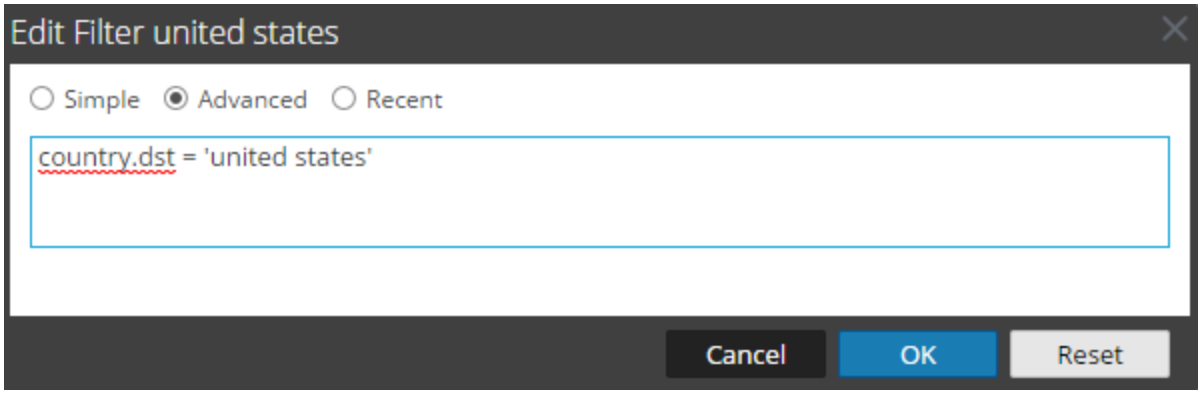
In the breadcrumb, you can click any of the crumbs to display the Query menu. You can delete a crumb and edit a query in a crumb. After each edit in the breadcrumb, NetWitness refreshes the results.

To work with queries in the breadcrumb:

1. Click a crumb.
   The Breadcrumb menu is displayed.



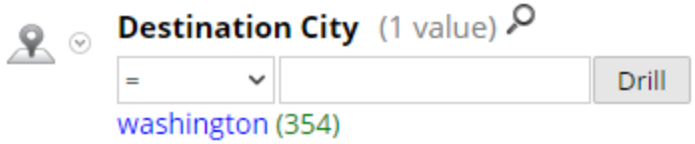
2. To edit a query in the breadcrumb, select **Edit**.
   The Create dialog is displayed with the selected query open for editing.



3. Edit the fields as described in Create a Query in the Navigate and Legacy Events Views.

To quick search within a meta key:

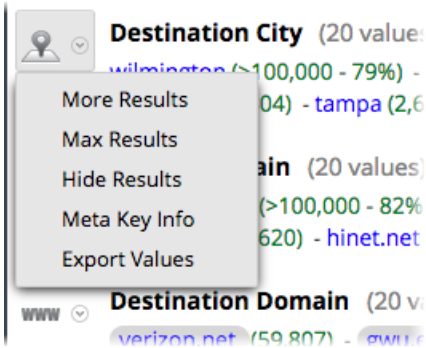
1. Move the mouse over a meta key section and click the magnifying glass.
   The Quick Search form, which contains a comparator and an optional operand for the search, is displayed.



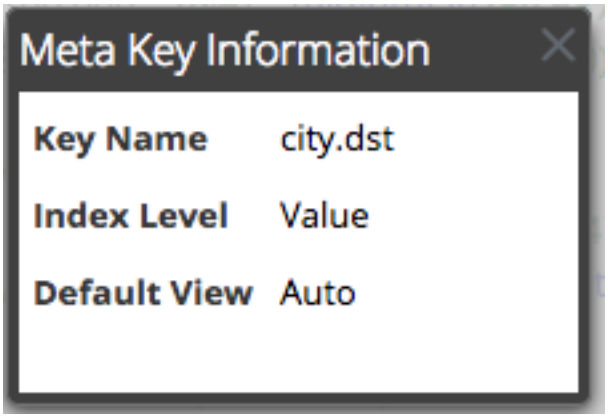
2. (Optional) If you want to close the search form, click the magnifying glass again.
3. Select the operation from the drop-down list on the left and type the text value to search for. Then click **Drill** to perform the execution.
   The metadata for that meta key is used to drill down in the current metadata.

To view meta key information and copy meta values for a meta key:

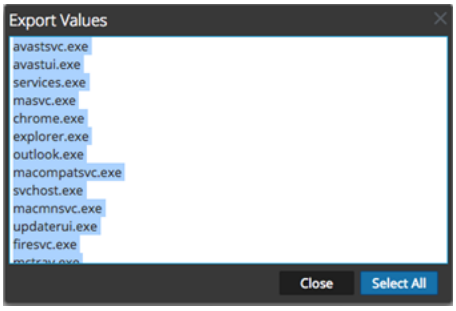
1. To view the key name, index level set for displaying the meta key, and the default view set for the meta key, click the drop-down menu next to the meta key. This figure shows the drop-down menu for Version 11.1 and later.



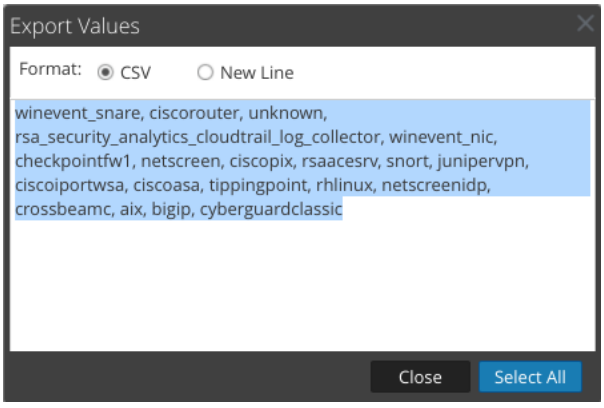
2. Select **Meta Key Info**.
   The Meta Key Info dialog is displayed.



3. When finished viewing, click ☒.
4. (Optional for Version 11.1 and later) To view meta values found for a meta key in a simple list that you can copy, click the drop-down menu next to the meta key.
   The Export Values dialog is displayed.
   The Version 11.1 dialog displays a list of values with one value per line.



   The Version 11.3 dialog allows you to select the method of separating values: either New Line or CSV.



5. Select the values that you want to copy, and click **Export Values**.
   The values are copied to the local clipboard and you can paste them into a file to save or share them.
6. To close the dialog, click **Close**.

7. (Optional) If you want to hide the results for the meta key in the current drill point, click the drop-down menu next to the meta key and click **Hide Results**.

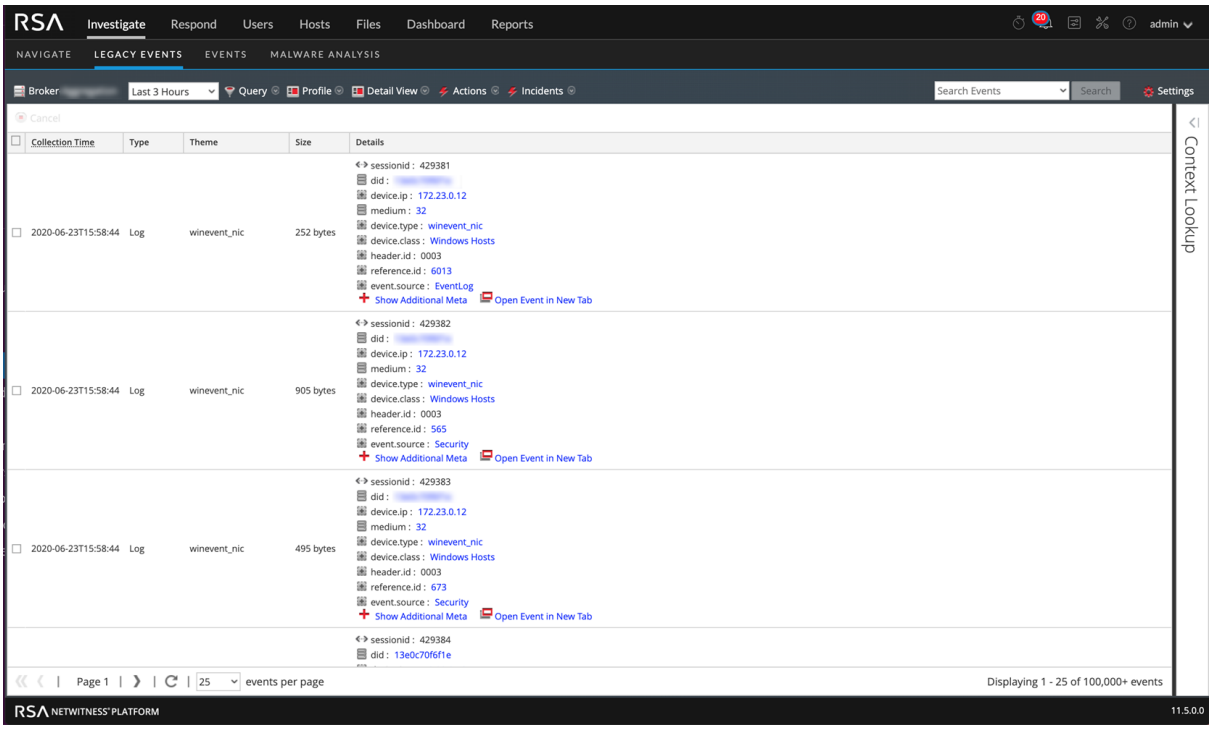
To display events associated with a meta value:

The Legacy Events view provides additional details for an event in two different views: Events List and Detail View.

1. In the Navigate view, drill into metadata that is the focus of your investigation.
2. Click the count (the number in green) next to a blue meta value.
   The Events view corresponding to the current drill point is displayed.
   The operations that you can perform in the events view are described in Reconstructing and Analyzing Events.

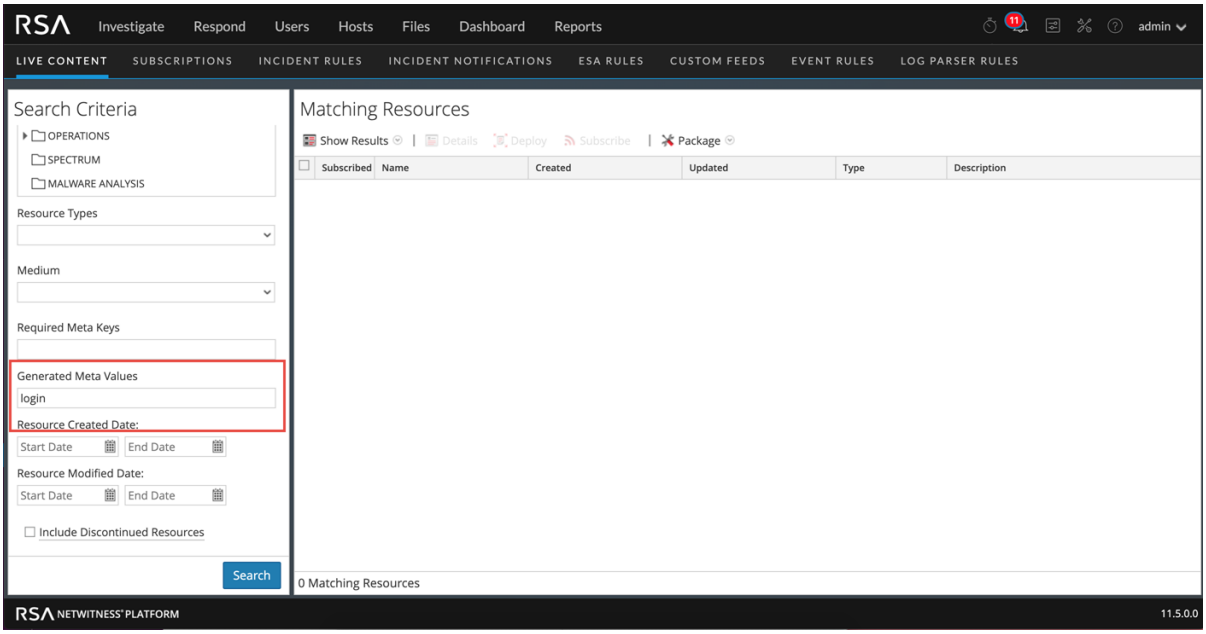
To search for specific events associated with a meta value:

1. In the **Navigate** view, drill into metadata that is the focus of your investigation (click a meta value or add a query).
2. Type a search string in the Search box and press **Enter** or click **Search**.
   You can also select and set search mode preferences. See Search for Text Patterns in the Navigate and Legacy Events Views for detailed search information.
   The Events view opens in a new tab and shows the search results. If you do not see the search term highlighted, click **Show Additional Meta**. Your time range selection and drills (queries) carry forward to the Events view.



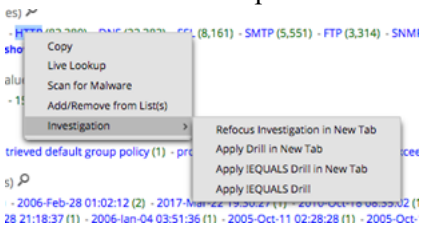
To view a selected meta value in RSA Live:

1. In the Navigate view, drill into metadata that is the focus of your investigation.
2. Right-click a meta value (the text in blue).
   The Meta Value drop-down menu is displayed.
3. To look up the meta value in RSA Live, select **Live Lookup**.
   The Live Search view is displayed with the meta value entered in the Generated Meta Value(s) field, and ready for a search.



To refocus the investigation in a drill point:

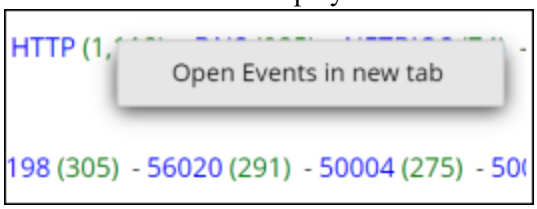
1. Right-click a meta value (the text in blue).
   The Meta Value drop-down menu is displayed.



2. Choose one of the refocus options.
   The drill is refocused according to your choice.

To look at a specific count in a new tab:

To view a count for a meta value in the Legacy Events view or the Events view, right-click a count for a meta value (the green number following the blue meta value).
The context menu is displayed.



Filter Results in the Legacy Events View

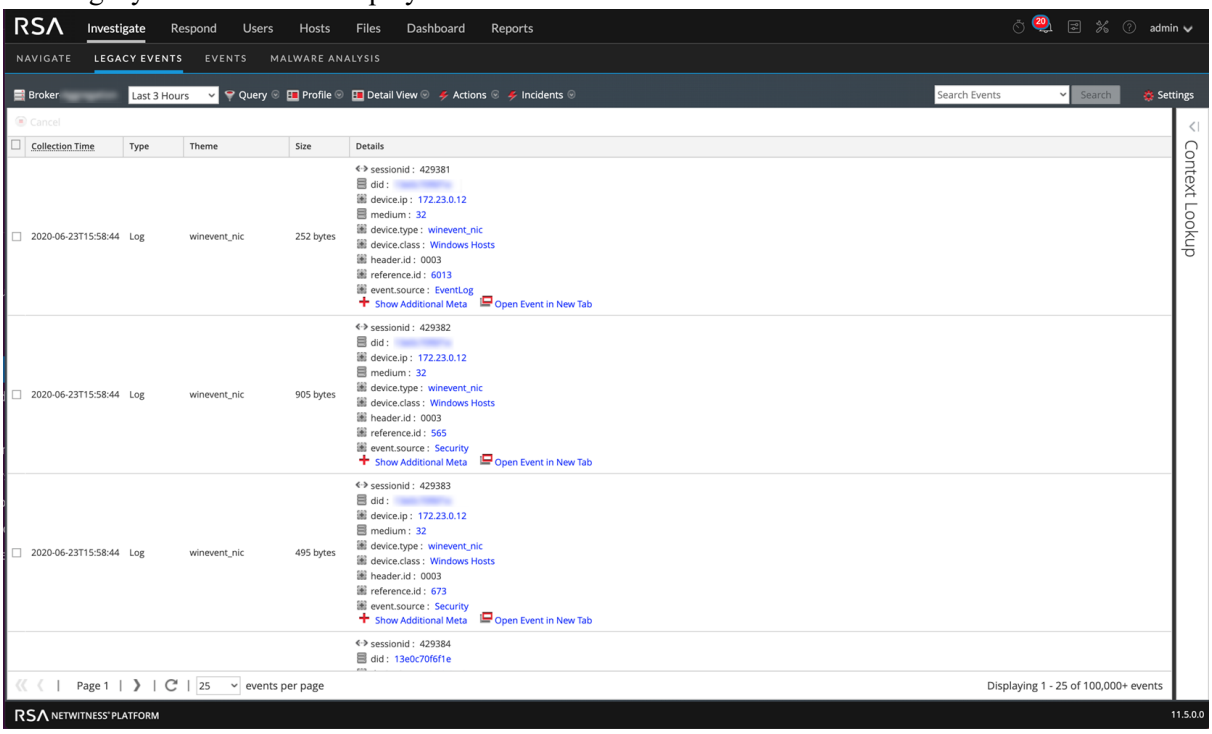
Analysts can filter events in the Legacy Events view by searching for events or selecting the service, setting the time range, and querying the metadata. If you opened the Legacy Events view from a Navigate view drill point, the view opens to the Detail view of events by default. Analysts who do not have permissions to use the Navigate view can query services directly from the Legacy Events view.

Note: When an Archiver is the currently selected service in the Legacy Events view and you are searching against a Broker or Concentrator, the search is slower than if searching against a Broker or Concentrator because the data on the Archiver is compressed and there is typically more data.

Filter Events Displayed in the Legacy Events View

To filter the data displayed in the Legacy Events view:

1. Go to **Investigate** > **Legacy Events**.
   The Legacy Events view is displayed.



2. To select a time range other than the default (**Last 3 Hours**), in the toolbar, click the time range field and select a value. For example, **Last Hour**.
   The Legacy Events view is refreshed with the selected time range.
3. Create a query as described in Create a Query in the Navigate and Legacy Events Views.
   The matching results for the query are displayed in the Detail View in the Legacy Events view. The breadcrumb reflects the query. In the breadcrumb, you can click any of the crumbs to display the

Query menu. You can insert a new query before a crumb, and append a new query to the end of breadcrumb. After each edit in the breadcrumb, the results are refreshed.

Page Through Events in the Legacy Events View

Pagination controls allow more flexibility in paging through a list of Events in the List View, Logs View, or Details View. You can select the number of events to display per page, and your selection is saved across logins to the NetWitness application. When a control is unavailable, the control is dimmed; for example, when you are viewing page 1, the ❮ and ❮❮ controls are dimmed.

To use pagination controls:

1. With results displayed in the Legacy Events view, click the current number of events per page (**10**, **25**, **50**, **100**, or **200**), and select the new number of events per page from the drop-down menu.
2. To page forward or back, use the page control icons:
   Click ❯ to go to the next page.
   Click ❯❯ to go to the last page.
   Click ❮ to go the previous page.
   Click❮❮ to go to the first page.
3. To go to a specific page, type a page number in the page number field | 3 Page 3 |.

Create a Query in the Navigate and Legacy Events Views

In the Navigate view or the Legacy Events view, you can create a query using dialogs for that offer syntax help with drop-down lists of applicable meta keys or meta entities and operators.

When viewing the drop-down list, you can expand and collapse each meta group to view or hide the individual meta keys in that group. When you select a meta group, NetWitness generates the complex query equal to a query with all of the meta keys in that group ORed together. So if a meta group contains `ip.src` and `ip.dst`, the query generated is `ip.src = <value> OR ip.dst = <value>`. If the meta group contains meta keys that have different meta value types, the value input is disabled and the query uses `exists` statements. For example, a meta group that contains `ip.src`, `ip.dst`, and `alias.host` includes meta keys that have different value types; `ip.src` and `ip.dst` are ip addresses and `alias.host` is text. The generated query is `ip.src exists OR ip.dst exists OR alias.host exists`.

A basic query is in the following form:

```
<metakey> <operator> [<metavalue>]
```

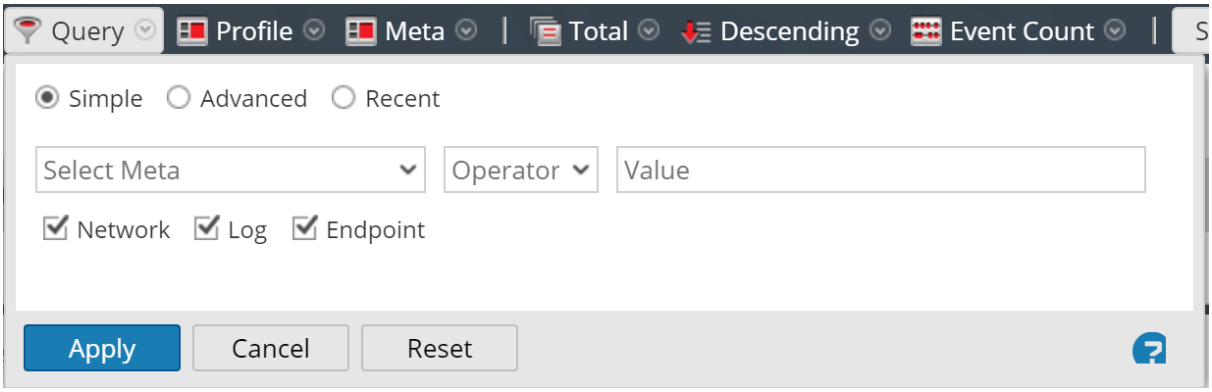
These are a few examples:

```
action exists
action = 'get'
alias.host = '10.25.55.115'
extension = 'exe'
orig_ip != "10.0.0.0" - "10.255.255.255"
```

Create a Query Using the Basic Method

When you create a query using the basic method, drop-down lists of meta keys and operators are displayed.

1. In the **Navigate** view or the **Legacy Events** view toolbar, select **Query.**
   The Query dialog is displayed, with the Simple option selected.

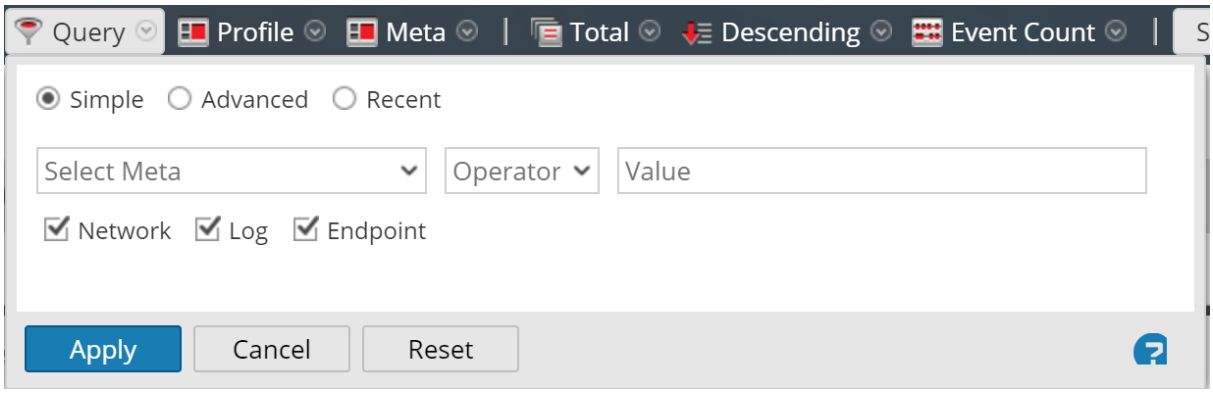


2. In the **Select Meta** field, click to display the drop-down list. The drop-down list has two sections: Meta Groups and All Meta.
3. Select a single meta key under **All Meta** or select a meta group under **Meta Groups**. You can also type in a meta key or meta group in the field.

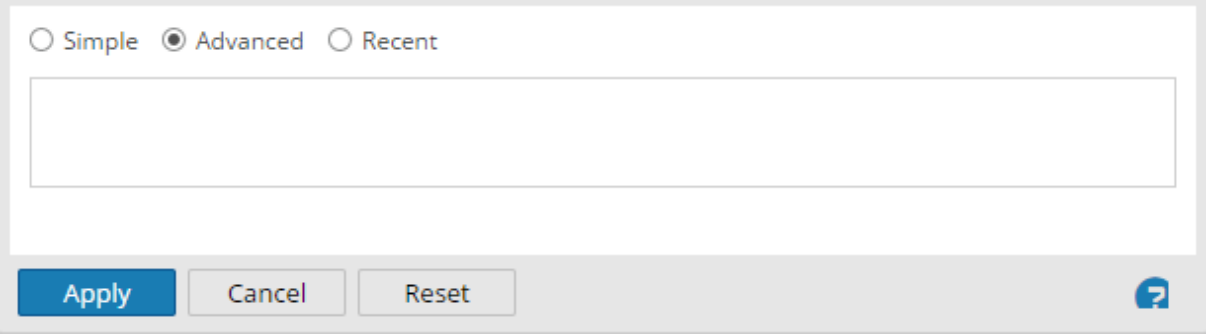
4. In the **Operator** field, type an operator or click on the drop-down list to select a valid operator.
5. (Optional) If you selected an operator that requires a value, for example, =, in the third field type the value for the meta key.
6. In the Network, Log, and Endpoint checkboxes, choose the type of data to query. Do one of the following:
   a. To limit the query to packets select **Network** and de-select **Log** and **Endpoint**.
   b. To limit the query to logs, select **Log** and de-select **Network** and **Endpoint**.
   c. To limit the query to endpoint events, select **Endpoint** and de-select **Network** and **Log**.
   d. To apply the query to packets, logs, and endpoints, select **Network**, **Log**, and **Endpoint**.
7. Do one of the following:
   a. Click **Apply**.
   The window is closed and the view is updated with the results of the new query. The query is displayed in the breadcrumb.
   b. Click **Cancel**.
   The window is closed and no changes are made to the view or current query.

Create a Query Using the Advanced Method

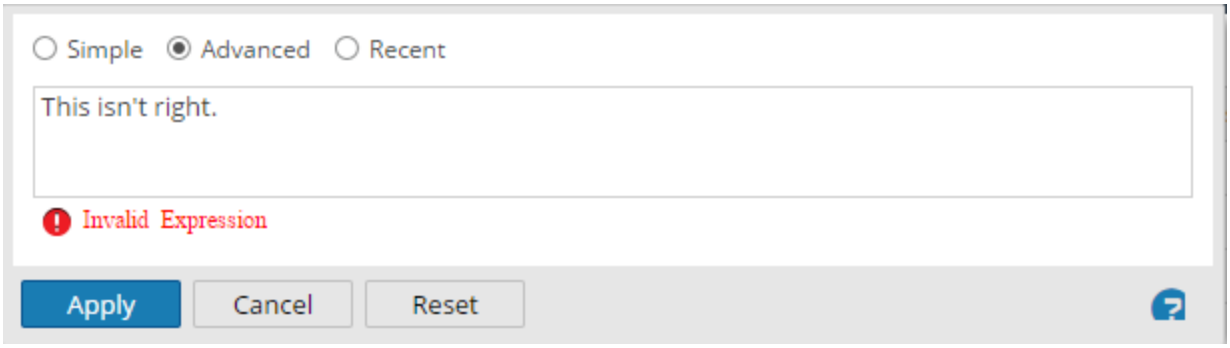
1. In the **Navigate** view or the **Legacy Events** view toolbar, select **Query.**
   The Query dialog is displayed.



2. Select **Advanced**.
   The advanced query field is displayed.



3. In the field, create a query, which can include the meta key, operator, and value. When you begin typing a meta key in the field a drop-down list of available meta keys for the selected service is displayed.
4. Select the meta key for your query.
   The display is updated. If the expression is not yet complete, the status indicates that the query is invalid.
5. Continue with an operator, from the drop-down list, then a value if necessary. The display is updated as you continue to enter the query. If you enter an operator, such as **exists** or **!exists**, which does not use the value field, the value field is disabled and the invalid status is cleared. If you enter an operator, such as =, which requires the value field, the invalid status remains until you enter a value. When the query is valid the invalid status is no longer displayed.

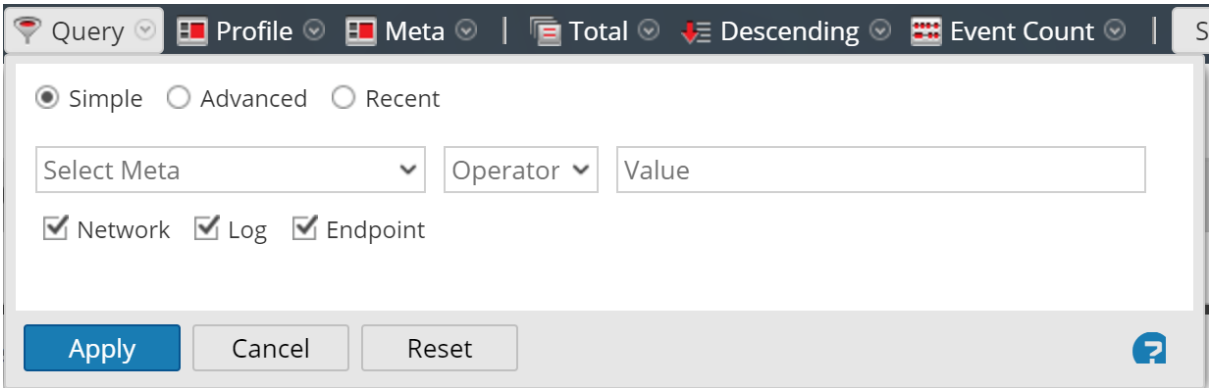


6. Do one of the following:
   - Click **Apply**.
     The window is closed and the view is updated with the results of the new query. The query is displayed in the breadcrumb.
   - Click **Cancel**.
     The window is closed and no changes are made to the view or current query.

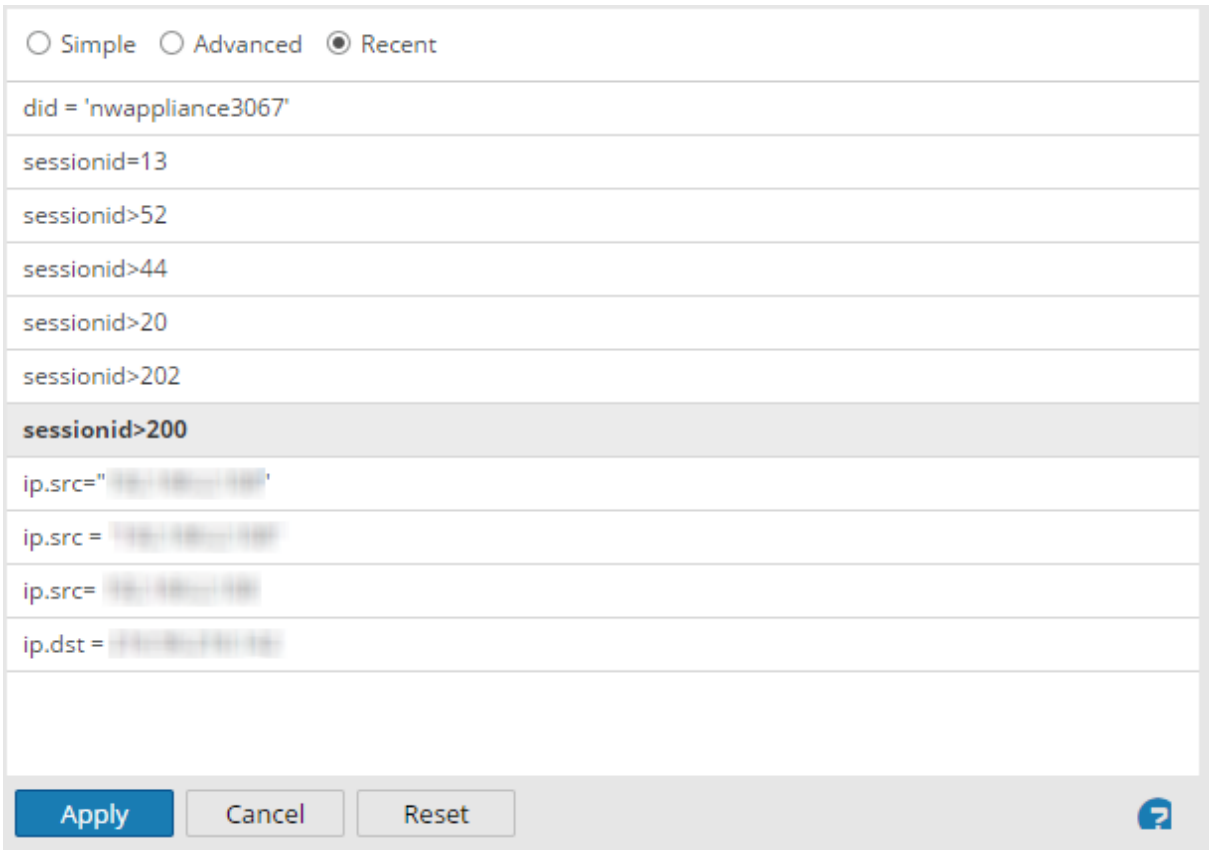
Apply a Recent Query

You can view recent queries and select one to apply to the current service being investigated. To select a recent query:

1. In the **Navigate view** or the Events view toolbar, select **Query.**
   The Query dialog is displayed, with the Simple option selected.



2. Select the **Recent** option.
   The list of recent queries is displayed in the bottom portion of the dialog.



3. In the list of recent queries, click to select a query.
4. Do one of the following:
   - Double-click a query.
   - Select a query and click **Apply**.
     The window is closed and the view is updated with the results of the new query. The query is displayed in the breadcrumb.
   - Click **Cancel**.
     The window is closed and no changes are made to the view or current query.

Search for Text Patterns in the Navigate and Legacy Events Views

You can search for text patterns within the current set of events in the Navigate view, the Events view, and the Legacy Events view. This section provides information about searching in the Navigate view and the Legacy Events view.

You can perform a keyword text search or do regex (Regular Expression) matching. In the Navigate view, you can click a meta value, such as HTTP, to drill into the data and then enter a search string in the Search field to search for events within that subset of data. The search opens a tab in the Legacy Events view, brings your drill and time range forward, and shows your search results. You can also drill into the data using queries before starting a search. To execute the search, enter a search string in the Search box, and press **Enter** or click **Search**.

Note: By default search results are only for exact matches found in indexed data. Only meta values shown as blue links in the Events Detail view are indexed. The regex option must be selected if the value contains a space. To broaden the search change the options in the Search Events drop-down menu.

Keyword Text Search

The text search provides these capabilities:

- Each white space delimited word is ANDed, so that every word must be found, but the order or location position in relation to the other words is irrelevant. For example, if you search on `Mark Albert`, both Mark and Albert must be found in the session, but they need not be together or in any specific order.
- The word OR is special. If you search `Mark OR Albert`, either Mark or Albert must be found in the session to match; both are not required.
- You can mix or match implicit ANDs and ORs together in the search string. The explicit OR has higher precedence than the implicit (whitespace) AND. The following examples make the same logical statement, which requires that both the terms cheese and dumplings be present in a match and one of toast or bread:

```
cheese toast OR bread dumplings
cheese AND (toast OR bread) AND dumplings
```

- You can exclude words from search results using the `-` operator. For example, searching for `cheese -toast` would return any result that has the word cheese, unless the word toast is also present.
- The keyword search can match metadata stored in the following patterns:
  - **IPv4 and IPv6 addresses**. Any term that can be recognized as an IP address is converted to the native metadata format so that it can be found in indexed metadata.
  - **IPv4 CIDR ranges**. You can use CIDR notation to locate IPv4 addresses within a range.

- **Timestamps**. Timestamps are matched against the native time metadata, and any additional time meta fields stored with the Time type.
- **Numbers**. The search function will attempt to automatically identify decimal search terms and match them against numeric meta data fields.

Options Controlling Search Behavior

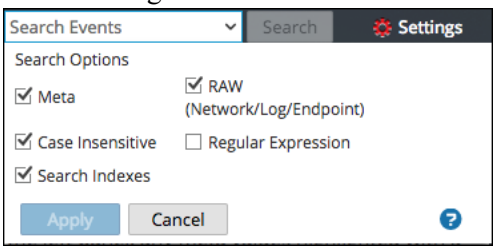
To access the Search box and search options in the Navigate or Legacy Events view:

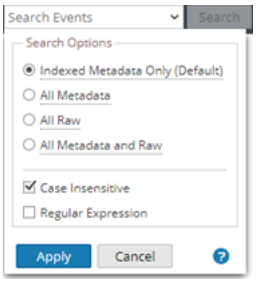
1. You can see the Search Events field in the toolbar.



Note: If you cannot see the Search Events field in the toolbar, click ⋯ on the right side of the toolbar.

2. Click in the **Search Events** field to view the Search Options drop-down menu. In Version 11.2 and later, the menu options are slightly different. The first figure illustrates the menu for 11.1 and below; the second figure illustrates the menu for Version 11.2 and above.





The options selected in this box change how the search is executed. The default search mode is to search indexed metadata and raw data only.

Note: Because the Index or Indexed Metadata Only (default) checkbox is selected by default, the search returns results based on data that is indexed. If you want to search for a complete set of metadata or raw data, select those checkboxes and clear the Index or Indexed Metadata Only (default) checkbox. This type of search takes longer, but it contains a more complete set of data.

The following table describes the Investigation search options.

Feature	Description
Indexed Metadata Only (default) checkbox (Version 11.2 and later) **Index** radio button (Version 11.1)	This search only returns results on indexed data. Searching the index is the fastest way to locate keywords within a large data set. The index search uses any relevant indexes present within your data collection. **Caution:** Substring matches are not located by index searches. If you require substring matches, clear this checkbox and use a non-index search mode.
All Metadata radio button (Version 11.2) **Meta** checkbox (Version 11.1)	Searches the metadata. Your keyword or regex pattern is matched against any parsed metadata.
All Raw radio button (Version 11.2 and later) **RAW** (Network/Log/Endpoint) checkbox (Version 11.1)	Searches the network, log, and endpoint event text. Every event is decoded and content is searched for matches on the keyword or regex pattern. If you select all data with no filters on an Archiver, execution time may be excessive and a warning may be displayed. **Caution:** Searching raw network sessions causes sessions to be decoded, which is very time intensive. You may want to disable raw searches when looking at network-only collections.
All Metadata and Raw radio button (Version 11.2)	Searches the metadata and the log or event text. This option is a combination of two options in Version 11.1: Meta and RAW (Network/Log/Endpoint), which you could select together. In Version 11.2, you can select only one radio button.
Case Insensitive	Ignores case when searching.
Regular Expression	Searches using a Perl regular expression, rather than text. By default executes a text search. To execute a regular expression search, select the Regular Expression option. **Caution:** - Regular expression searches can be very slow. - When combining regular expressions and index search options, the regular expression pattern is matched against unique index values instead of meta values. This produces results faster, but it is not an exhaustive search of all the metadata or raw data.
Apply	Sets the default search options to apply to a search in the Navigate and Legacy Events views. This also updates your Investigation preferences in your Profile (Profile > Preferences > Investigation tab). The preferences are saved and effective immediately. You can select search options to use for a particular search without changing your default search preferences.

Regular Expression Search Syntax

A regular expression search uses Perl regular expression syntax, which is documented in detail in http://perldoc.perl.org/perlre.html.

Raw Text Keyword Search

The Log Decoder has the capability to create a raw text index for unparsed log events. This functionality creates metadata items that form a full-text index on downstream services such as Concentrators and Archivers. When you enable the Search Indexes option in your search preferences, your search automatically uses the text index. Note that the text index produces meta items that have a coarse granularity. For example, the default text indexer configuration truncates text terms. By comparing the index matches against raw data, the search engine will find accurate results for your search. However, you can improve search times by disabling the raw search checkbox. If you do so, results will be returned faster, but you may see false positive hits in your search results.

Search Procedures

Search in the Navigate View

To search within the currently displayed data in the Navigate view:

1. Type a search string in the Search field and press **Enter** or click **Search**.
2. To clear the search box and return to the previous Navigate view with results unfiltered by the search, click the **X** in the search box.

Search in the Legacy Events View

To search within the currently displayed data in the Legacy Events view:

1. Type a search string in the Search box, and press **Enter** or click **Search**.
   The search results are displayed. Events that match the search criteria are displayed in the events list. In the Details view and List view, matches are highlighted in the Details column. In addition, when searching RAW, matches are highlighted in the Log view Logs column.
2. If you want to narrow the search, change the query and time.
3. If you want to stop the search and return to the Legacy Events view, click **Cancel**.
   Any results that are displayed remain.
4. To clear the search box and return to the normal Events view, click **X** in the search box.

View and Modify Queries Using URL Integration

NetWitness Investigate includes an External URL Integration that facilitates integration with third-party products by allowing a search against the NetWitness architecture. By using a query in a URI, you can pivot directly from any product that allows custom links, into a specific drill point in the Investigate view. This integration provides an internal presentation of the user's query.

URL Integration allows the user to identify the service either by the host id or by the service and port, as defined in NetWitness. If NetWitness is unable to resolve the service, the analyst is redirected to the Navigate view, showing the Service selection dialog. Once the service is selected, the Navigate view is loaded with the drill point, defined by the query.

Service Id Known

When the ID of the service to use for an investigation is known, the format for entering a URI using a URL-encoded query is:

```
http://<sa host:port>/investigation/<deviceId>/navigate/query/<encoded query>/date/<start date>/<enddate>
```

where

- `<sa host: port>` is the IP address or DNS, with or without a port, as appropriate (ssl or not). This designation is needed only if access is configured over a non-standard port through a proxy.
- `<deviceId>` is the internal Service ID in the NetWitness instance for the service to query against. The service ID can be represented only as an integer. You can see the relevant service ID from the URL when accessing the Investigate view within NetWitness. This value changes based on the service being connected to for analysis.
- `<encoded query>` is the URL-encoded NetWitness query. The length of query is limited by the HTML URL limitations.
- `<start date>` and `<end date>` define the date range for the query. The format is `<yyyy-mm-dd>T<hh:mm:ss>Z`. The start and end dates are required. If no date is provided then the user defaults for that service are used. Relative ranges (for example, Last Hour) are not supported. All times are run as UTC.
  For example:
  `http://localhost:9191/investigation/12/navigate/query/alias%20exists/date/2012-09-01T00:00:00Z/2012-10-31T00:00:00Z`

Host and Port Known

When the host and port of the service to use for investigation is known, the format for entering a URI using a URL-encoded query is:

```
http://<sa host:port>/investigation/<device host:port>/navigate/query/<encoded query>/date/<start date>/<enddate>
```

where

- <sa host: port> is the IP address or DNS, with or without a port, as appropriate (ssl or not). This designation is needed only if access is configured over a non-standard port through a proxy.
- <device host:port> is the host and port of a service defined in the NetWitness instance for the service to query against. NetWitness attempts to resolve the host and port as a service ID defined in NetWitness.
- <encoded query> is the URL-encoded NetWitness query. The length of query is limited by the HTML URL limitations.
- <start date> and <end date> define the date range for the query. The format is <yyyy-mm-dd>T<hh:mm:ss>Z. The start and end dates are required. If no date is provided then the user defaults for that service are used. Relative ranges (for example, Last Hour) are not supported in this version. All times are run as UTC.
  For example:
  `http://localhost:9191/investigation/concentrator:50105/navigate/query/alias%20exists/date/2012-09-01T00:00:00Z/2012-10-31T00:00:00Z`

Examples

These are query examples where the NetWitness Server is 192.168.1.10 and the deviceID is identified as 2.

All activity on 03/12/2013 between 5:00 and 6:00 AM with a hostname registered

- Custom Pivot: `alias.host exists`
- `https://192.168.1.10/investigation/2/navigate/query/alias%2Ehost%20exists/date/2013-03-12T05:00:00Z/2013-03-12T06:00:00Z`

All activity on 3/12/2013 between 5:00 and 5:10 PM with http traffic to and from IP address 10.10.10.3

- Custom Pivot: `service=80 && (ip.src=10.10.10.3 || ip.dst=10.0.3.3)`
- Encoded Pivot Dissected:
  - `service=80 => service&3D80`
  - `ip.src=10.10.10.3 => ip%2Esrc%3D10%2E10%2E10%2E3`
  - `ip.dst=10.10.10.3 => ip%2Esrc%3D10%2E10%2E10%2E3`
  - `https://192.168.1.10/investigation/2/navigate/query/service%3D80%20%26%26%20%28ip%2Esrc%3D10%2E10%2E10%2E3%20%7C%7C%20ip%2Edst%3D10%2E10%2E10%2E3%29/date/2013-03-12T17:00:00Z/2013-03-12T17:10:00Z`

Additional Notes

Some values may not need to be encoded as part of the query. For example, commonly the IP src and dst is used for this integration point. If leveraging a third-party application for integration of this feature, it is possible to reference those without encoding applied.

Reconstructing and Analyzing Events

Having refined events in the Navigate view or in the Events list as described in Refining the Results Set, your next step is to learn more about the events by reconstructing them, looking at attachments, and viewing additional context in third-party lookups or internal lookups.

Reconstructions are done in the Events view or the Legacy Events view. If you are starting from the Navigate view, you need to go to the Events view or the Legacy Events view to see a reconstruction.

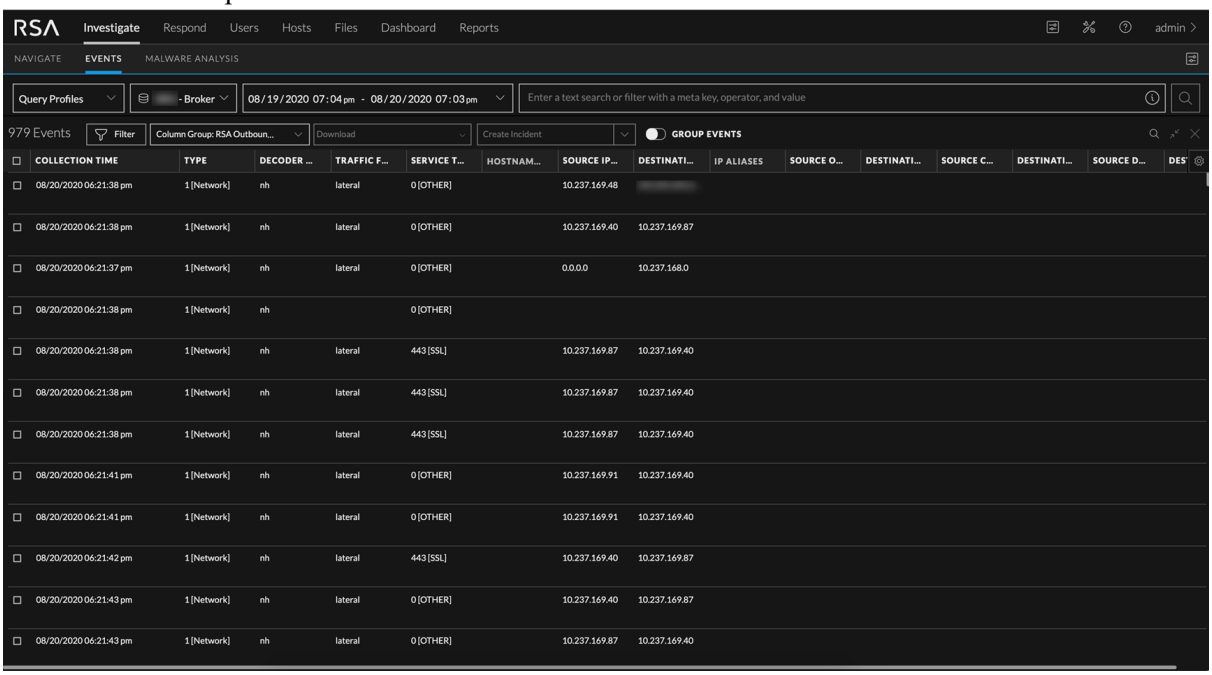
Note: The Legacy Events view is disabled by default. The administrator can enable the view as described in "Configure Investigation Settings" in the *System Configuration Guide*.

To display events in the Events view, do one of the following:

1. Go to **Investigate** > **Events**.
2. Go to **Investigate > Navigate** (Version 11.5 and earlier), right-click the meta count for a meta value (the meta count is in green text). When the context menu is displayed, select **Open Events in new tab**.



The Events view opens with a list of events for the selected meta value.

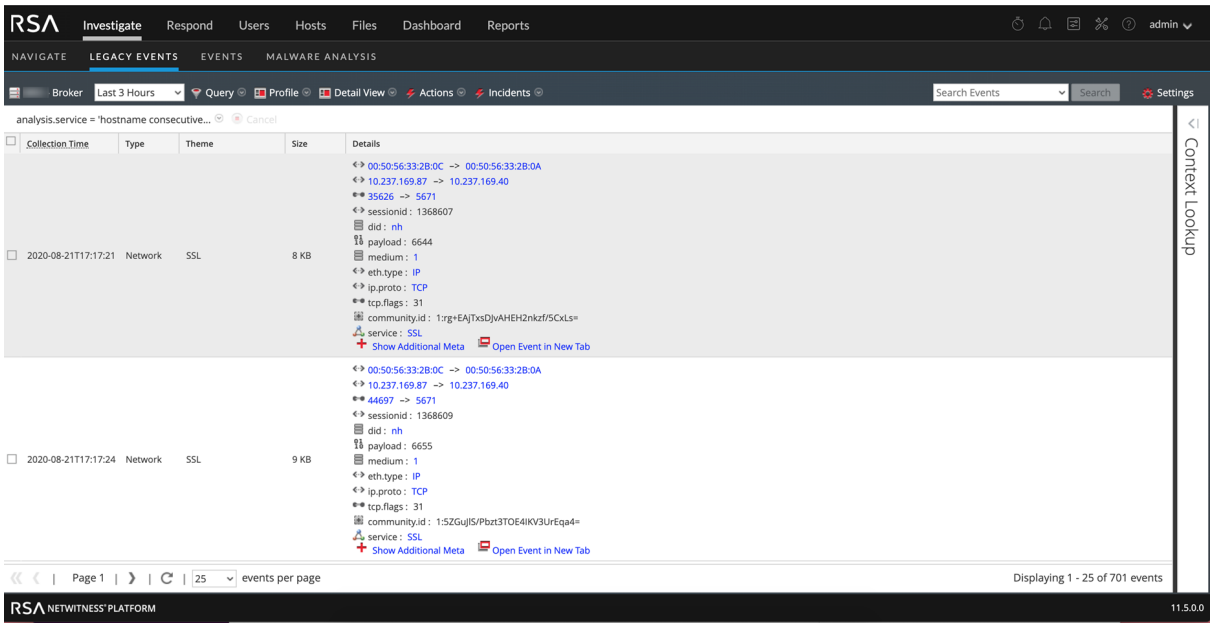


For detailed information about the types of reconstruction and analysis that you can use in this view, see Examine Event Details in the Events View.

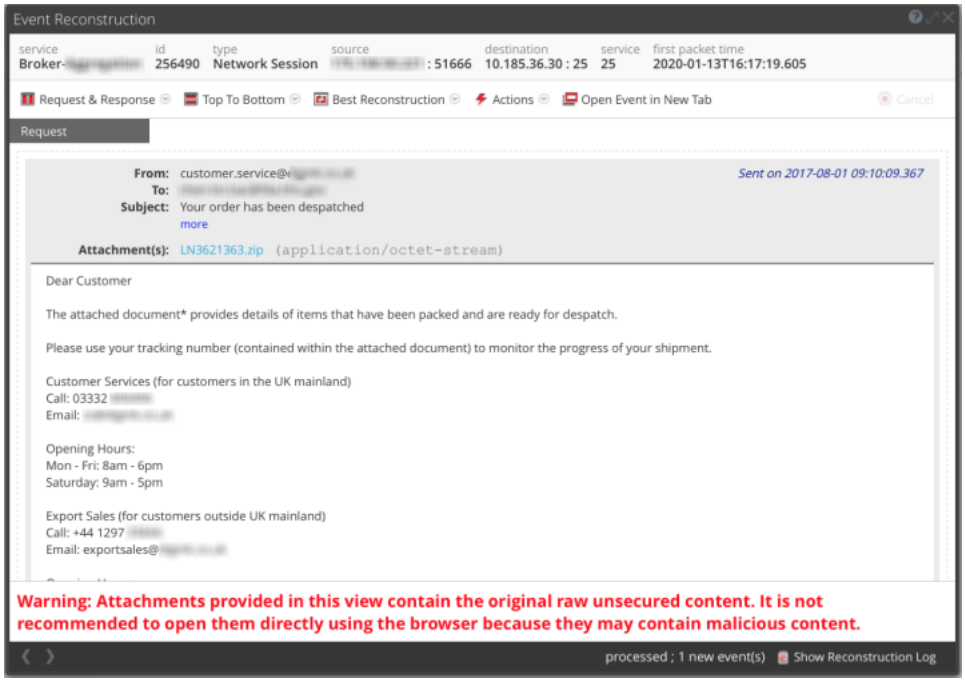
To display an event in the Legacy Events view, do one of the following:

1. To open the Legacy Events view using the default query for the default service, go to **Investigate** > **Legacy Events**.(This option is available only if the administrator has enabled the view.)
2. To view events for a specific meta value in the Legacy Events view, go to **Investigate > Navigate** and when events are loaded in the Values panel, click a meta count (the meta count is in green text). You can also right-click the meta count for a meta value. When the context menu is displayed, click **Open Legacy Events in new tab**.
   The Legacy Events view displays the events for the selected meta value. The Legacy Events view provides three built-in presentations of event data: the Detail view, the List view, and the Log view. This figure is an example of the Detail view. You can use queries, the time range setting, and profiles to filter the events listed in the Legacy Events view. You can extract files, export events, export logs, and open the Event Reconstruction panel by double-clicking an event. See Downloading and Acting Upon Results for detailed information about these capabilities.
   NetWitness runs a default query on the last three hours for the default service (if one is set) or displays a dialog in which you can select a service and then runs the default query. The default query selects all events and the Events view displays events on the selected service, with the oldest events

first.



3. To view a reconstruction of the first event in the list, double-click the event.
The reconstruction opens in a pop-up window in front of the Events list.



Examine Event Details in the Events View

When you find an interesting session in the Navigate view or the Events view > Filter Events panel, you can see the list of sequential events for the session in the Events view > Events panel. Clicking an event in the list opens the Network Event Details panel for that type of event: Network Event Details, Log Event Details, or Endpoint Event Details. Within the Event Details panel, you can select a tab that shows an event reconstruction (text, packet, file, email, and web) or (Version 11.5 or later) the tab that shows host information for network events that are enriched with endpoint data.

Note: (Version 11.5 or Later) For expanded network visibility of existing network events in your network (packet) deployment. Network events are enriched with endpoint data namely the host and process that triggered the network event and other details such as user name, risk score, reputation, and so on.
You can view endpoint data in the following ways:
- (Quick View) **Investigate** > **Events** - Event Summary Header
- (Detailed View) **Investigate** > **Events** > **Host**

For more information to enable expanded network visibility, see "Creating Groups and Policies" in the *Endpoint Configuration Guide*.

Note: For expanded network visibility to work, ensure the service user account used for aggregating Endpoint Log Decoder data to Endpoint Concentrator is assigned with the `decoder.manage` permission. For more information on how to assign roles and permissions, see "Add a Role and Assign Permissions" in the *System Security and User Management Guide*.

Event Details for Each Event Type

Within the Event Details panel, different tabs are available per event type as shown in the following table. Procedures for working in the Event Details panel are provided in Analyze Events in the Events View.

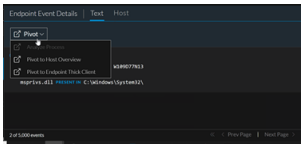
Action	Network Event	Log Event	Endpoint Event
View the text reconstruction (default unless last selected overrides)	✓	✓	✓
View the file reconstruction	✓		
(Version 11.5 or Later) View the host Information for an Endpoint Agent configured with expanded network visibility (see Host Information.	✓		
View the packet reconstruction	✓		
View the email reconstruction	✓		
View the web reconstruction in the Legacy Events view (see Reconstruct an Event in the Legacy Events View)	✓		

Each tab has settings to enhance your analysis. If you change a setting, the setting is preserved between browser refreshes and logins within the same browser. These are the preserved settings:

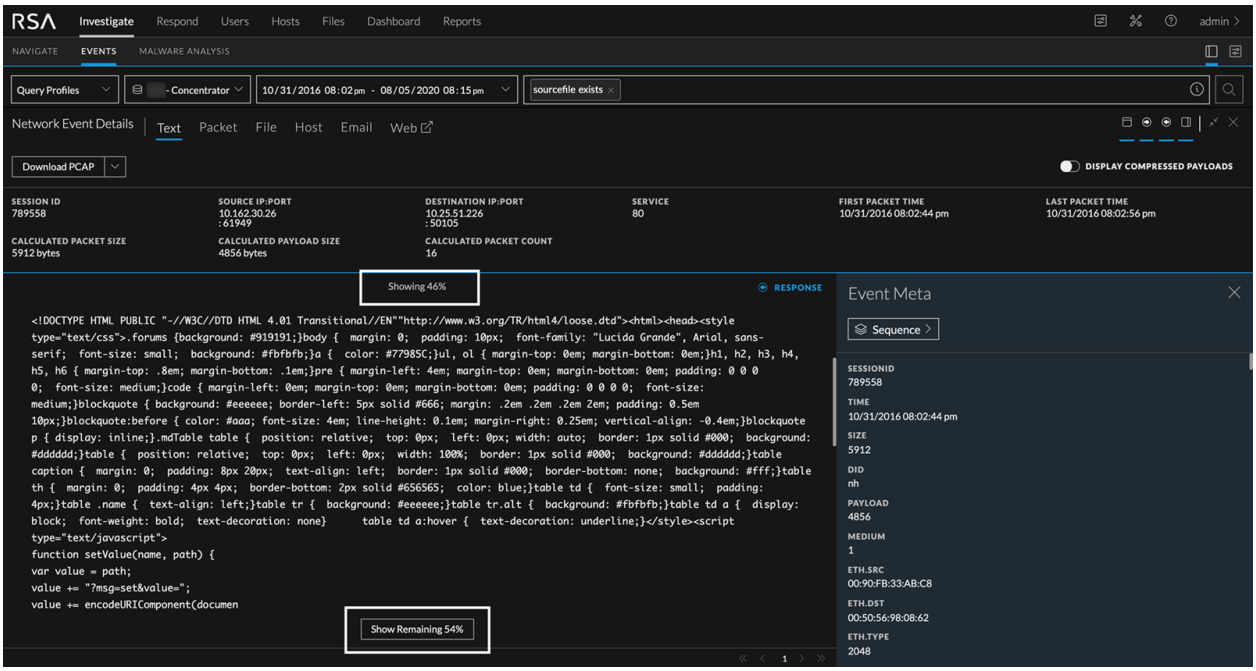
- The currently selected reconstruction: Text, Packet, File, (Version 11.5 and later) **Host**, or Email.
- Whether the Event Meta panel is open or closed.
- Whether the Event header is open or closed.
- Whether the Request or Response, or both are displayed.
- Whether packet payloads are displayed without the headers in the packet reconstruction.
- Whether shaded bytes are displayed in the packet reconstruction.
- Whether other common file types are highlighted in the packet reconstruction.
- The number of packets per page in the packet reconstruction.
- Whether compressed or uncompressed text is displayed in the text reconstruction.

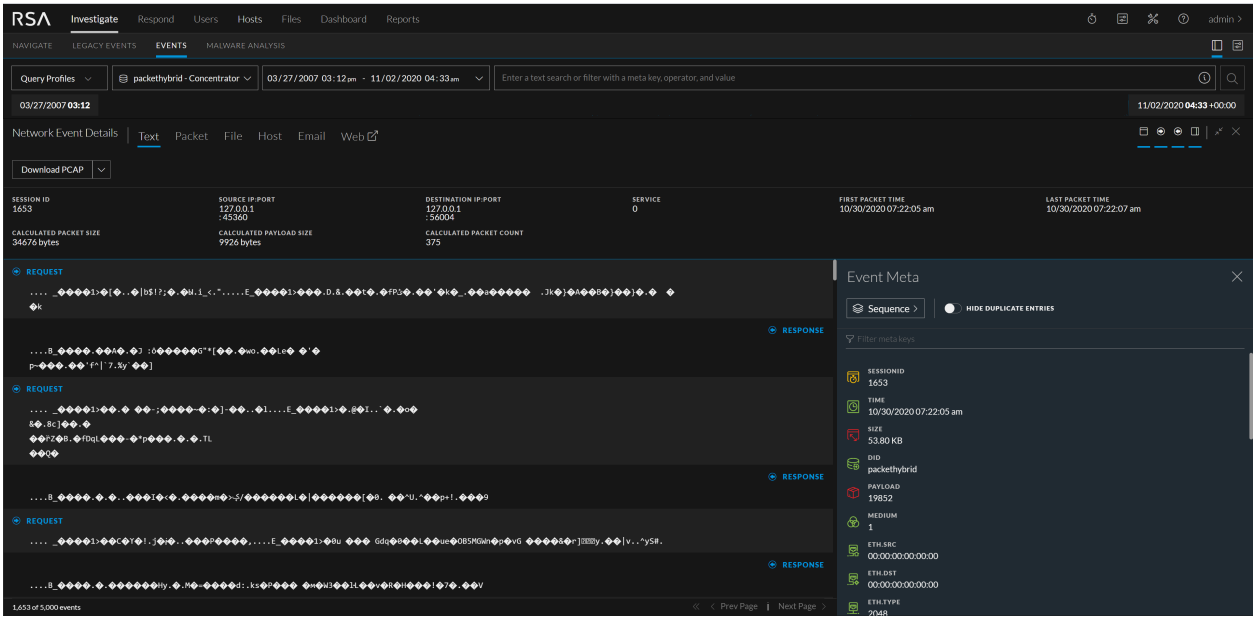
Text Reconstruction

For Endpoint Events Details, the Pivot options available under Text analysis page has been replaced with a Pivot dropdown menu that contains the three options to perform further investigation.



You can view all types of events (network events, log events, and endpoint events) in their original text format in the Text tab. The text reconstruction for some network events can be quite large. To ensure the best rendering, an excessively large payload is truncated to fit. If a single reconstructed request or response in the reconstructed event exceeds the maximum number of bytes, the header indicates what percentage of bytes is shown. Pagination controls add flexibility when paging through the reconstructed text of an event. This figure illustrates a single response that has been truncated because it exceeds the maximum number of bytes (Version 11.2 and later).





Note: Version 11.1 handles large payloads differently; the payload for a single event is limited to 2500 packets. When the packet limit is reached, a warning in the footer advises the limit has been reached and provides the total number of packets in the event. For Version 11.1, the Show More option is still available for messages that are truncated; however, the entire text of the message is not visible without downloading the raw payload.

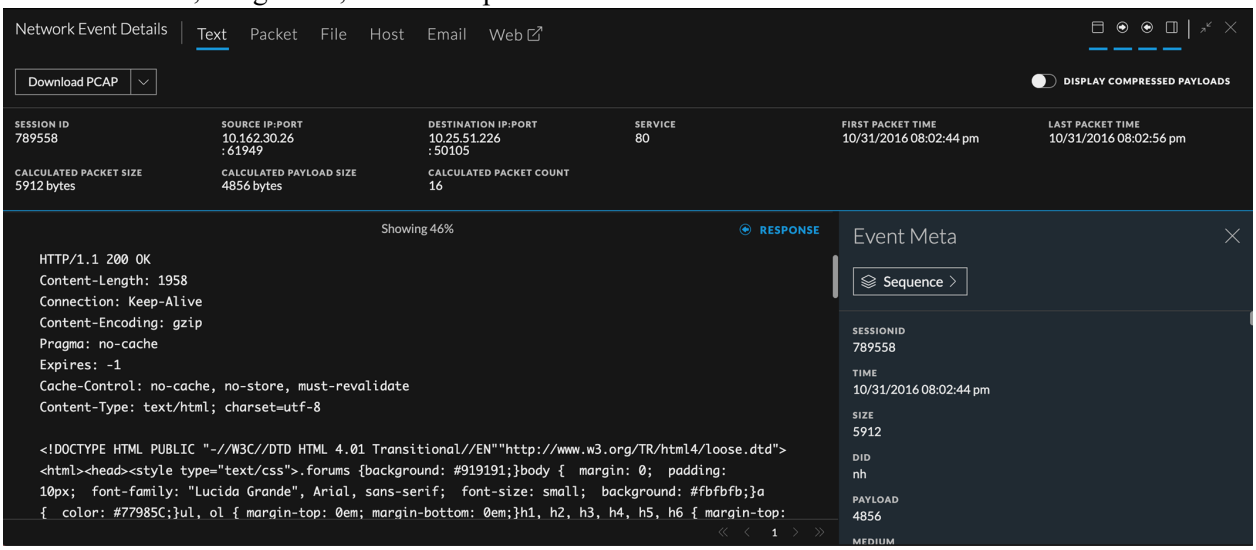
In the text reconstruction, network events, log events, and endpoint events are presented differently.

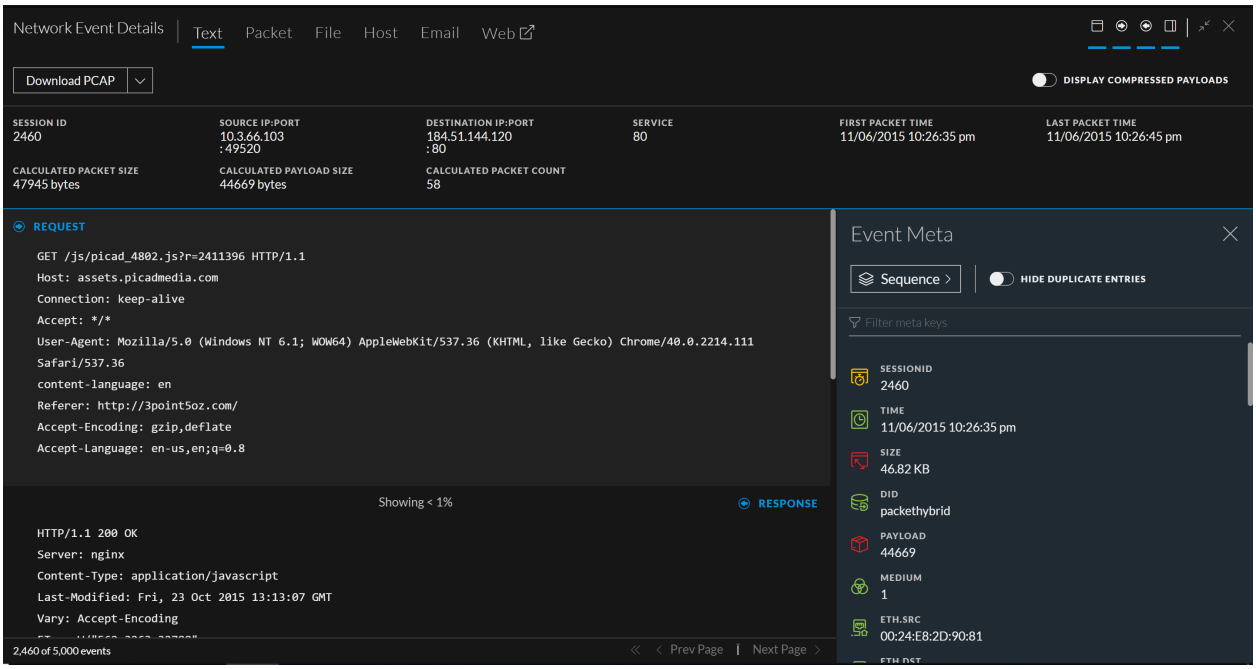
- For network events, the reconstruction provides the direction of the packet (Request or Response) and contents of each packet in text format. If you are reconstructing a network event, the text

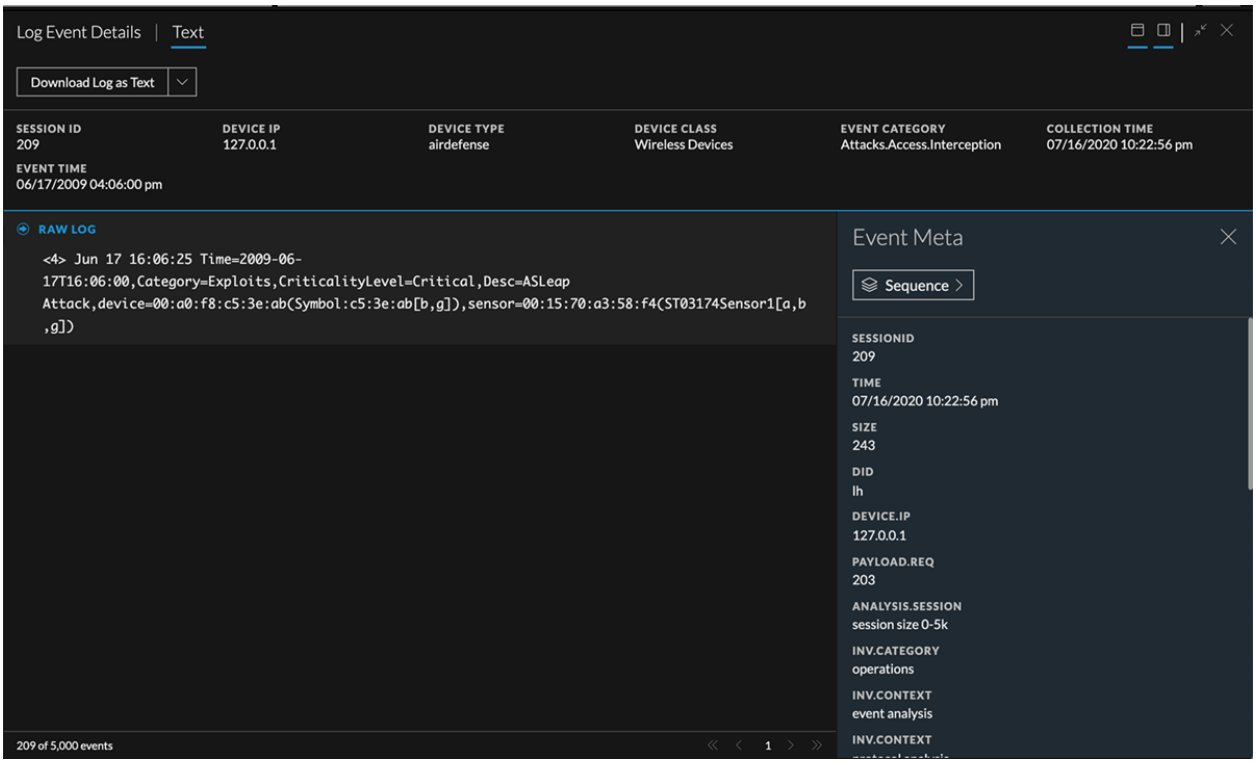
reconstruction is scrollable. When you scroll, the text identification information and the Request and Response labels remain visible rather than scrolling out of view.

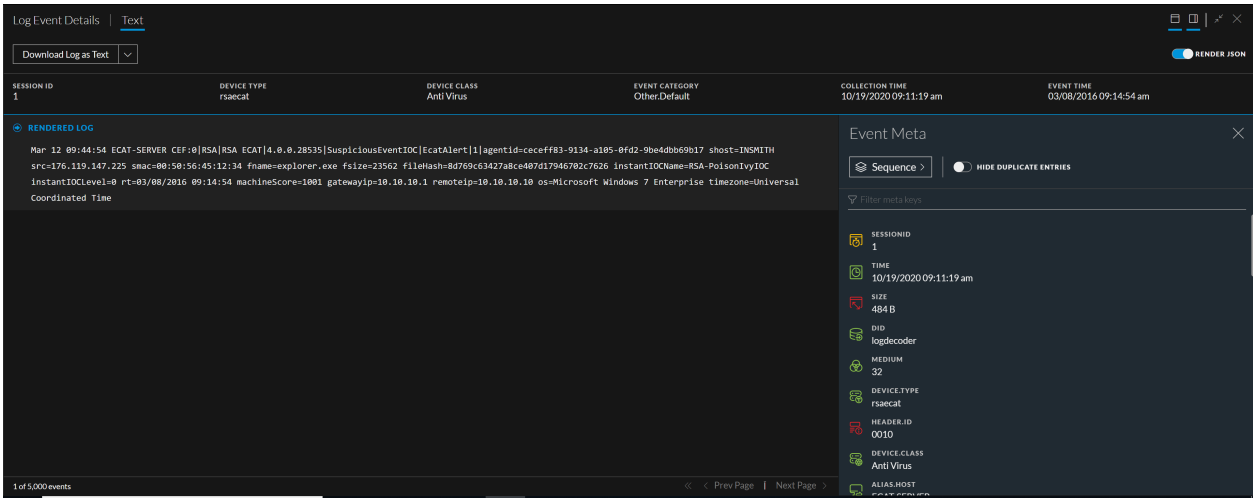
- Log events and endpoint events have no request or response; only the raw event is displayed in the Text tab. Endpoint events include additional information relevant to an endpoint event.
- (Version 11.5.1) Log events that contain JSON snippets are rendered in an easily readable JSON tree view with nested indentations if the RenderJSON option is enabled.

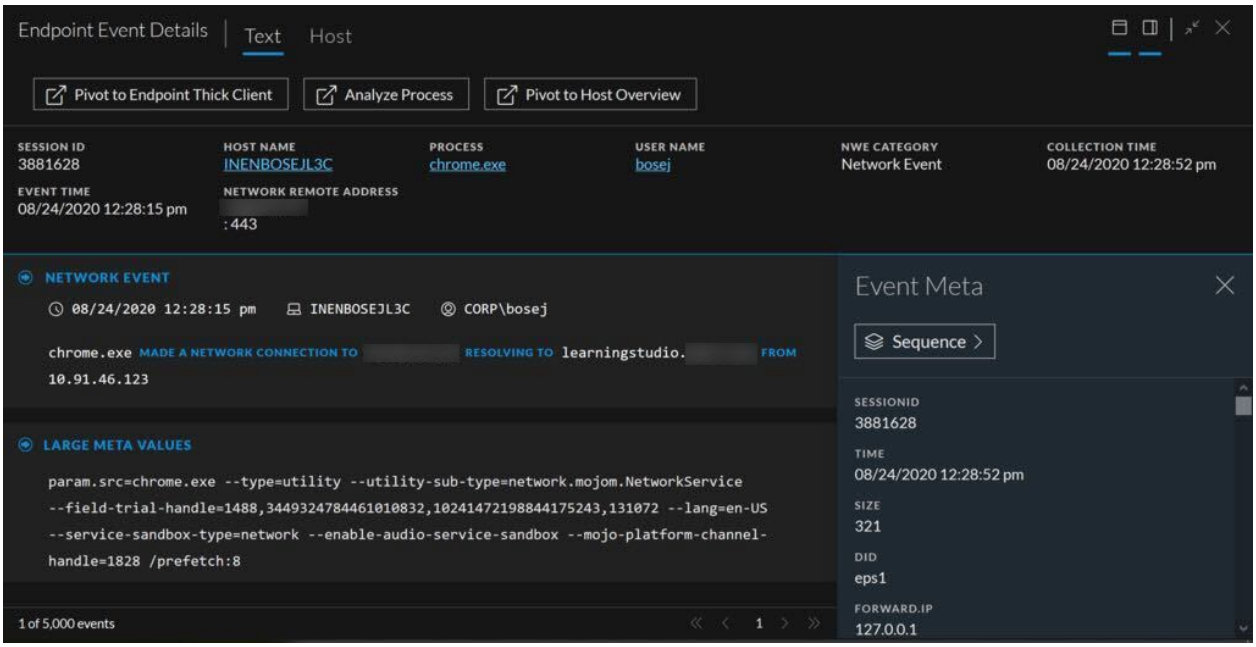
For each type of event (network, log, or endpoint), there are differences in the event header and the options for downloading the event. Below is an example of the text reconstruction for each type of event: a network event, a log event, and an endpoint event.






Log Event Details | Text
Download Log as Text
SESSION ID
209
DEVICE IP
127.0.0.1
DEVICE TYPE
airdefense
DEVICE CLASS
Wireless Devices
EVENT CATEGORY
Attacks.Access.Interception
COLLECTION TIME
07/16/2020 10:22:56 pm
EVENT TIME
06/17/2009 04:06:00 pm
RAW LOG
<4> Jun 17 16:06:25 Time=2009-06-17T16:06:00,Category=Exploits,CriticalityLevel=Critical,Desc=ASLeap Attack,device=00:a0:f8:c5:3e:ab(Symbol:c5:3e:ab[b,g]),sensor=00:15:70:a3:58:f4(ST03174Sensor1[a,b,g])
209 of 5,000 events
Event Meta
Sequence
SESSIONID
209
TIME
07/16/2020 10:22:56 pm
SIZE
243
DID
lh
DEVICE.IP
127.0.0.1
PAYLOAD.REQ
203
ANALYSIS.SESSION
session size 0-5k
INV.CATEGORY
operations
INV.CONTEXT
event analysis
INV.CONTEXT


Log Event Details | Text
Download Log as Text
RENDER JSON
SESSION ID
1
DEVICE TYPE
rsaecat
DEVICE CLASS
Anti Virus
EVENT CATEGORY
Other.Default
COLLECTION TIME
10/19/2020 09:11:19 am
EVENT TIME
03/08/2016 09:14:54 am
RENDERED LOG
Mar 12 09:44:54 ECAT-SERVER CEF:0|RSA|RSA ECAT|4.0.0.28535|SuspiciousEventIOC|EcatAlert|1|agentid=cecef83-9134-a105-0fd2-9be4dbb69b17 shost=INSMITH src=176.119.147.225 smac=00:50:56:45:12:34 fname=explorer.exe fsize=23562 fileHash=8d769c63427a8ce407d17946702c7626 instantIOCName=RSA-PoisonIvyIOC instantIOCLevel=0 rt=03/08/2016 09:14:54 machineScore=1001 gatewayip=10.10.10.1 remoteip=10.10.10.10 os=Microsoft Windows 7 Enterprise timezone=Universal Coordinated Time
1 of 5,000 events
Prev Page
Next Page
Event Meta
Sequence
HIDE DUPLICATE ENTRIES
Filter meta keys
SESSIONID
1
TIME
10/19/2020 09:11:19 am
SIZE
484 B
DID
logdecoder
MEDIUM
32
DEVICE.TYPE
rsaecat
HEADER.ID
0010
DEVICE.CLASS
Anti Virus
ALIAS.HOST

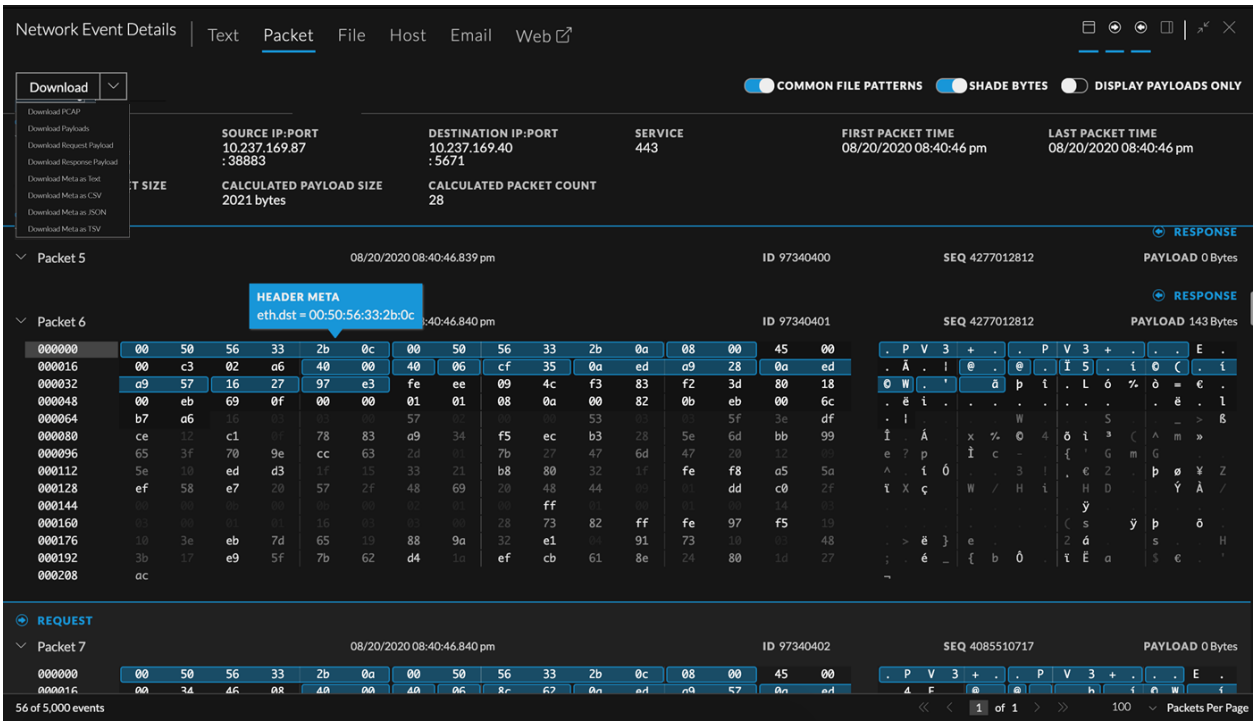


Note: The calculated packet count, calculated packet size, and calculated payload size in the Event header may be different than the same statistics in the Event Meta panel because the metadata is sometimes written before event parsing completes and may include packet duplicates.

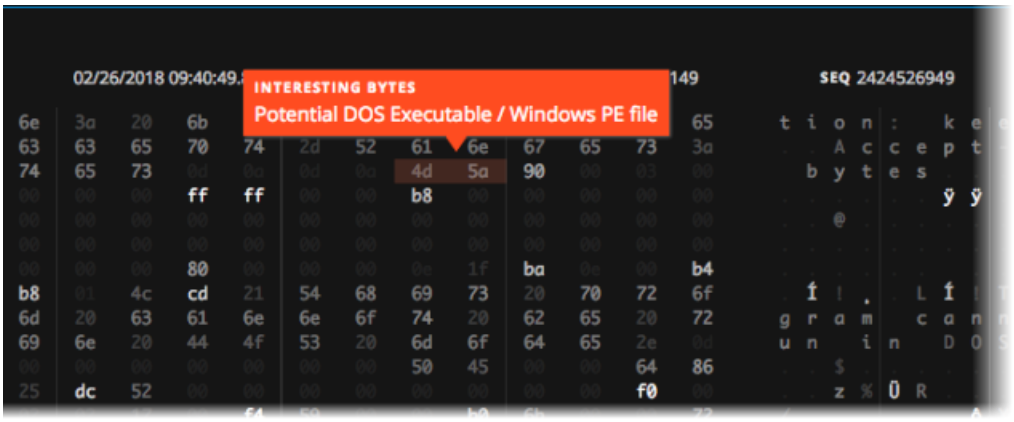
Packet Reconstruction

The packet reconstruction is for network events. The panel is scrollable, and the packet identification information and the Request and Response labels remain visible rather than scrolling out of view. In the Packet tab, the headings provide the direction of the packet (Request or Response), the packet number, the packet start time, the packet ID and the sequence, and the payload size. All packets begin with a header, and some packets have a footer. Pagination controls add flexibility when paging through packets.

The metadata in the hexadecimal and ASCII data is highlighted in blue; when you place the cursor over the highlighted metadata, the meta key/meta value information is displayed in a hover box.

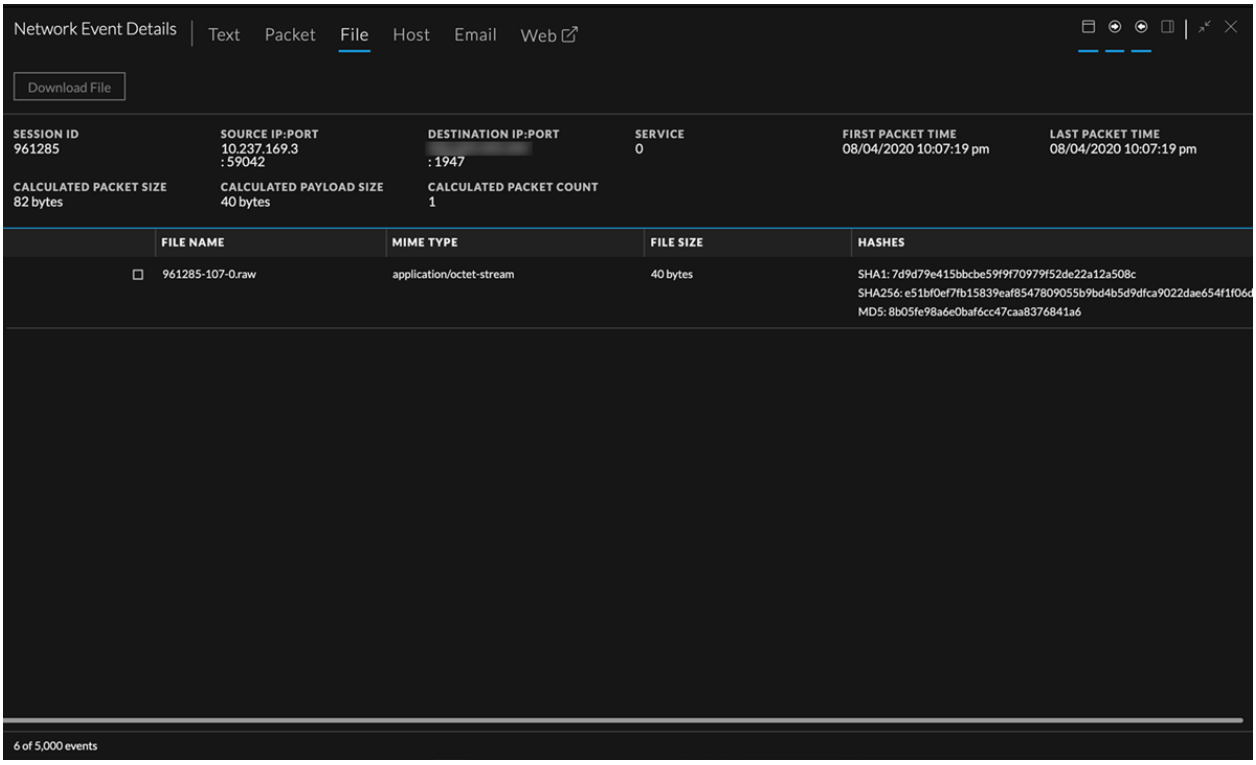


Common file signatures are highlighted with an orange background. When you place the cursor over the highlighted text, the description of the file type is displayed in a hover box.

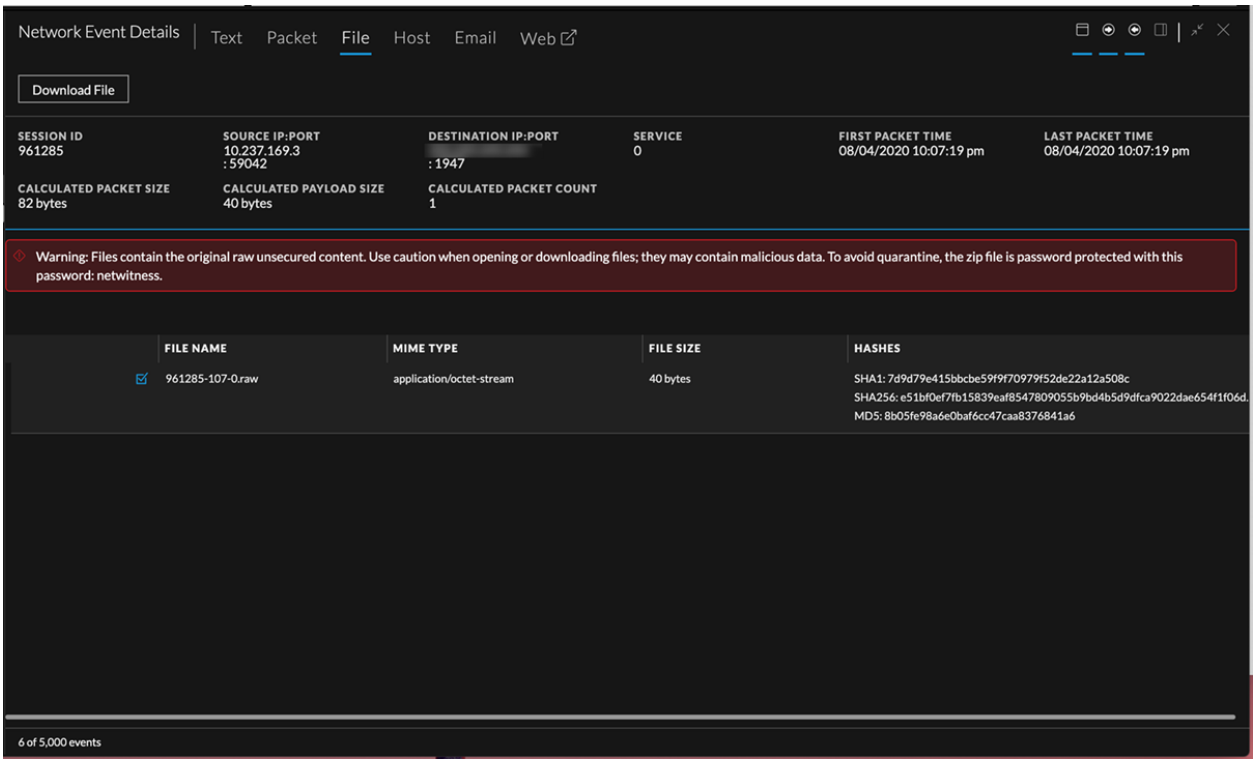


File Reconstruction

The file reconstruction shows a list of files associated with the selected network event. This is an example of the file reconstruction.



You can select one or more files, or all files, to export to your local file system. When files are selected, The Download File options becomes active and reflects the number of files selected.



Caution: Caution is advised when unzipping and opening files that are associated with a default application; for example, an Excel spreadsheet may automatically open in Excel before you have a chance to verify it is safe.

Host Information

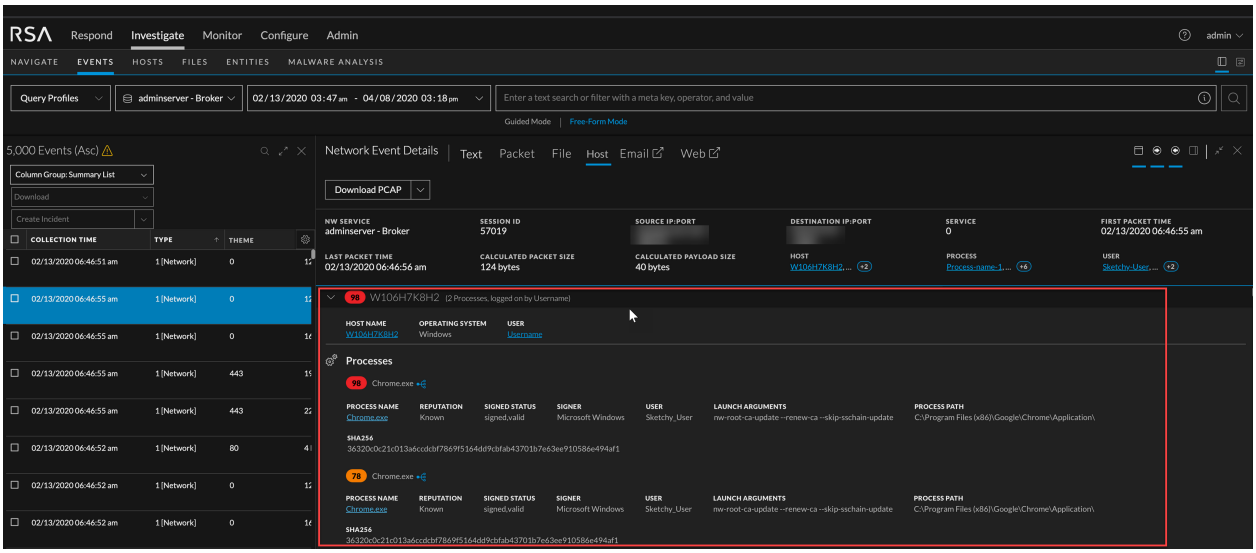
Host information is displayed for network and endpoint events with the endpoint data.

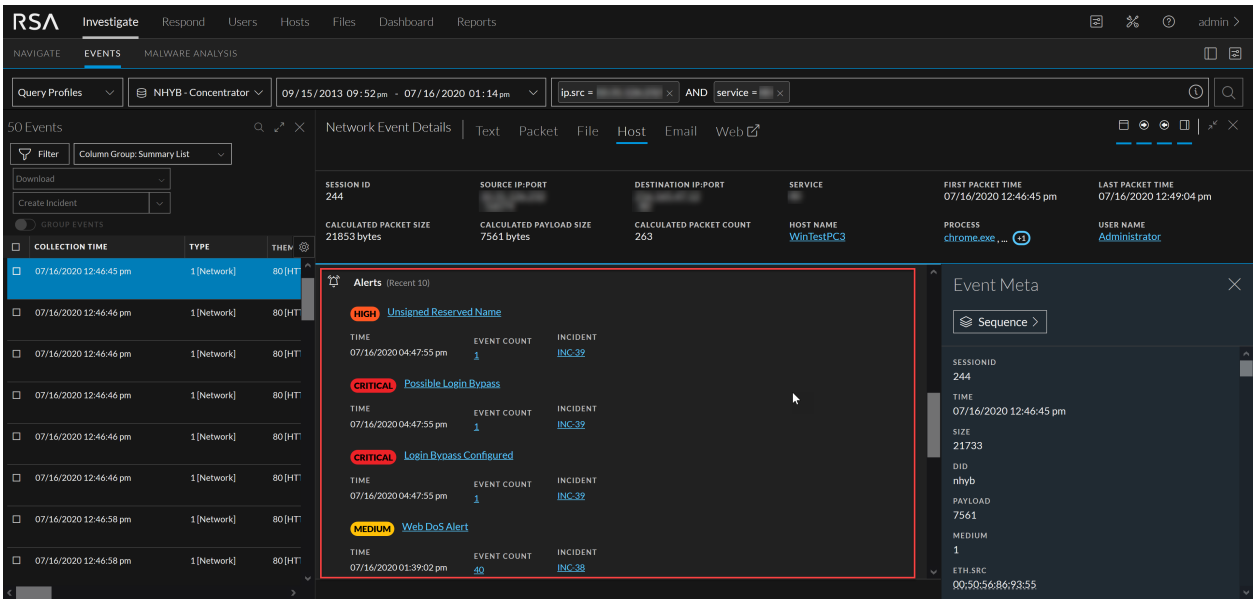
You can filter the events in the Events view using the Events Filter panel (version 11.5 and later). For more information, see Filter Results in the Events View.

Note: Endpoint data is displayed only if you have an Endpoint deployment, and the Endpoint agents are configured for expanded network visibility. For more information to enable expanded network visibility, see "Creating Groups and Policies" in the *Endpoint Configuration Guide*.

Note: For expanded network visibility to work, ensure the service user account used for aggregating Endpoint Log Decoder data to Endpoint Concentrator is assigned with the `decoder.manage` permission. For more information on how to assign roles and permissions, see "Add a Role and Assign Permissions" in the *System Security and User Management Guide*.

Below are some examples of the host information.





- The event summary header includes the following information from endpoint data:
  - Host name - The host from where the event originated.
  - Process - The source process that triggered the event.
  - User - The user associated with the triggered process.
- The hosts with the closest matching process are listed in the order of event time.
- By default, the first host is expanded, and you can view additional information such as:
  - Host details – This provides details on the host's operating system and the owner (logged in user) associated with the host.

- To investigate on the host name, click the **Host name** link highlighted in blue. For more information, see "Investigating Hosts" in the *NetWitness Endpoint User Guide*.
- To investigate alerts associated with the user, click the **owner** link highlighted in blue. For more information, see "Investigate High-Risk Entities" in the *NetWitness UEBA User Guide*

◦ Process details – This provides details like risk score, process name, reputation, event time, on hosts, signed status, process ID, signer, user, launch arguments, SHA256, and path.

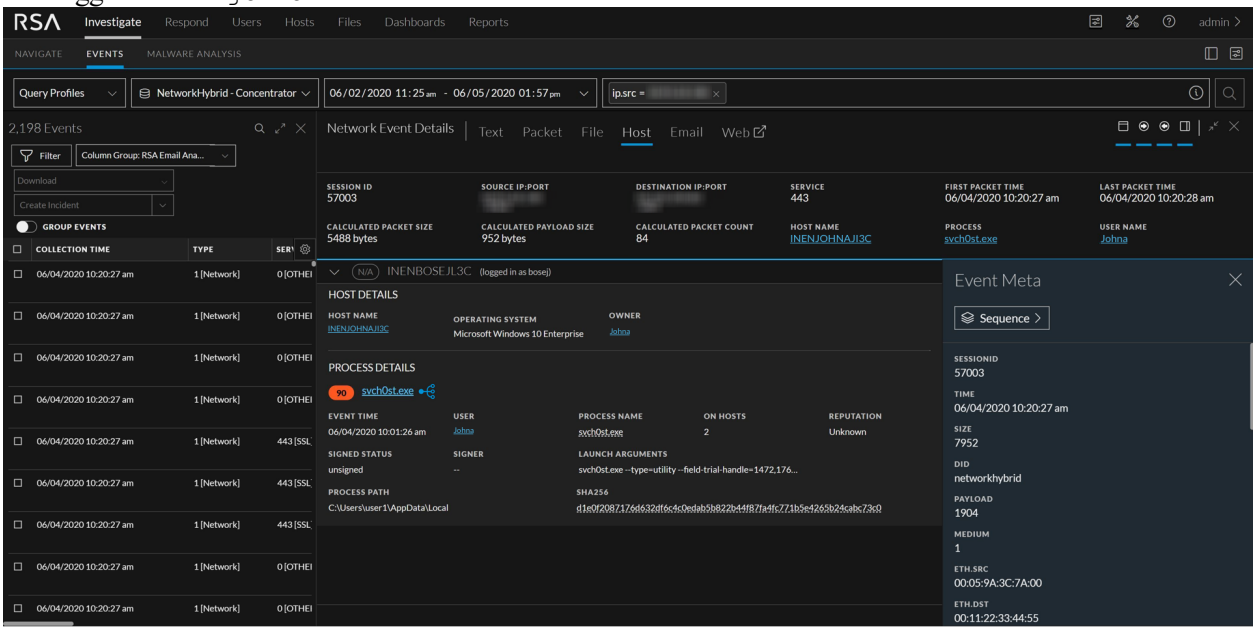
- Click  to open the process tree. By default, the process tree will open the process details of last 14 days. The icon to open the process tree does not appear when the process tree is not available.
- To investigate on the process, click the **process** link highlighted in blue. For more information, see "Investigating Files" in the *NetWitness Endpoint User Guide*.
- To investigate alerts associated with the user, click the **user** link highlighted in blue. For more information, see "Investigate High-Risk Entities" in the *NetWitness UEBA User Guide*.

◦ Alert details – It displays the recent ten alerts associated with the host. These alerts can be from endpoint, network, and log events. You can click **View All** to open the host details page. The host details page lists all the alerts that contribute to the risk score. You can click on the alert name to open the alert details. For information on how to review an alert, see “Reviewing Alerts” in the *Net Witness Respond User Guide*. This section provides the following details.

- Severity – Displays the severity of the alert.
- Time – Date and time when the alert was triggered.
- Event Count – It displays the number of events that triggered the alert. To view the events associated with the alert, click the **EVENT COUNT** link highlighted in blue. The **EVENT COUNT** link is available only when the events are from the same source.
- Incident – Lists the incidents associated with the respective alerts. To view the details and respond to an incident, click the **INCIDENT** link highlighted in blue. For more information, see "Responding to Incidents" in the *NetWitness Respond User Guide*.

You can hover over the meta values of the host name, process, user, owner, and SHA256 to view additional information about the specific metadata. For more information on context look up, see Look Up Additional Context for Results.

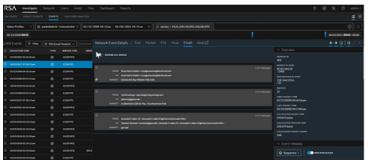
Below is an example of the Host Information tab with a single host, process, and user associated with the selected network event. The `svch0st.exe` is the process associated with the host `INENJOHNAJI3C` and logged in user `johna`.



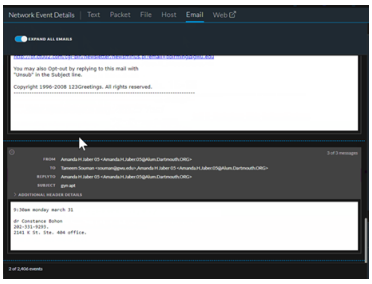
Note: You may see multiple hosts and processes triggered for the selected network event; in such cases, the host from where the event is triggered first is listed first and then the other hosts where a similar event is triggered.
For example, if `10.63.0.240` IP address is assigned to Host1, and User1 is logged in to the machine and accessed `www.nyu.edu/` using Chrome. Meanwhile Host1 is powered off (within a span of 30 minutes), and the same IP address is assigned to Host2. The user logged in is User2 and accesses `www.nyu.edu/` using Internet Explorer. In this case, network events for the endpoint data are as follows:
- Hostname - `Host1`, `Host2`
- Process - `chrome.exe`, `iexplore.exe`
- User - `User1`, `User2`

Email Reconstruction

In 11.7 and later, if the analysts need to review all email contents in a single session, then they can click on the Expand All Emails toggle button by navigating to **Investigate** > **Events** > **Email** view.

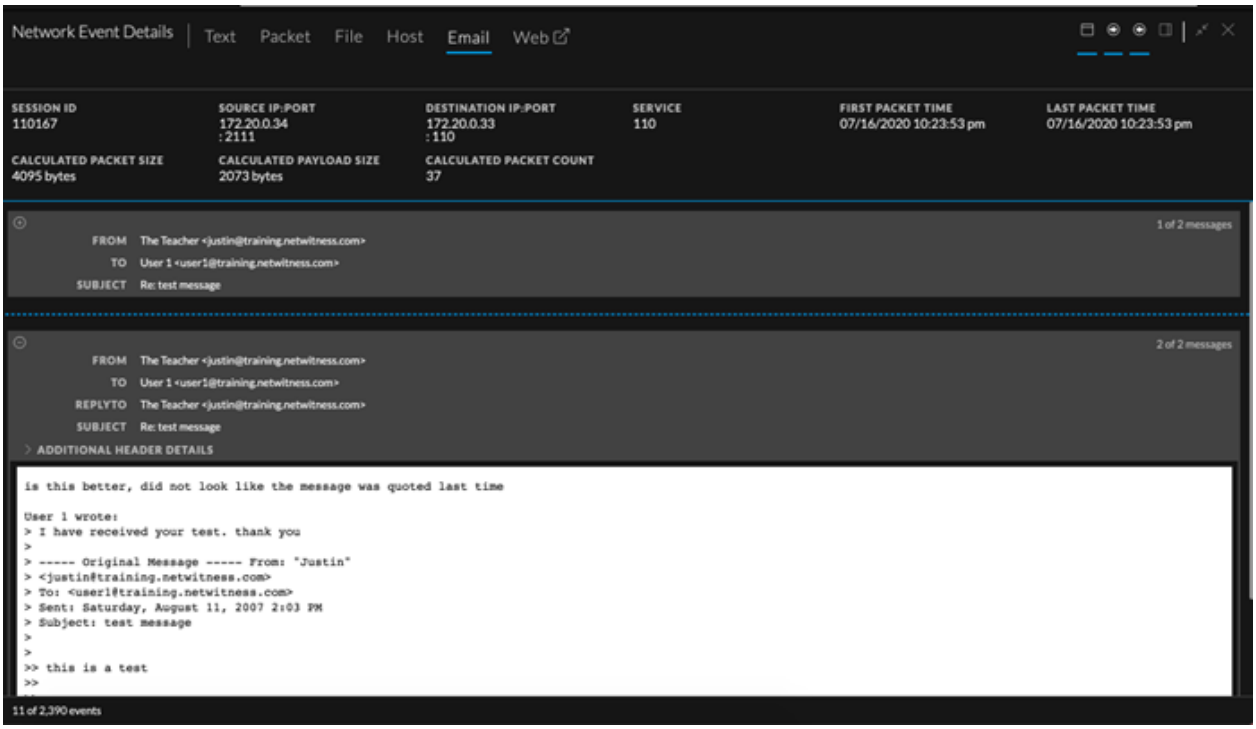


When the Expand All Emails toggle button switched on, the email content is displayed in an expanded form.



When the Expand All Emails toggle button is switched off, the email content is displayed in a collapsed form. If there are no emails to display the toggle button is disabled.

The email reconstruction shows a list of emails associated with the selected network event. This is an example of the email reconstruction.



- By default, a single email is expanded and multiple emails are collapsed.
- If an email contains attachments, you can download attachments as described in Download Data in the Events View.

  Caution: When you download and open attachments from an email, they may contain malicious data.
- An external link in an email cannot be accessed. Clicking an external link displays a **Link Address** popup window that provides the actual link.
- When an email body is too long, **Showing %** is displayed in the beginning of the email. To view the remaining content, click **Show Remaining %** at the bottom of the email.
- If an event contains a web email supported by the `alias.host` metadata of `mail.google.com`, `mail.live.com`, or `mail.yahoo.com`, a message is displayed with a link to view the

reconstruction for the associated session in the Event Reconstruction page. If not, a "*No Email reconstruction is available for this event*" message is displayed.

Analyze Events in the Events View

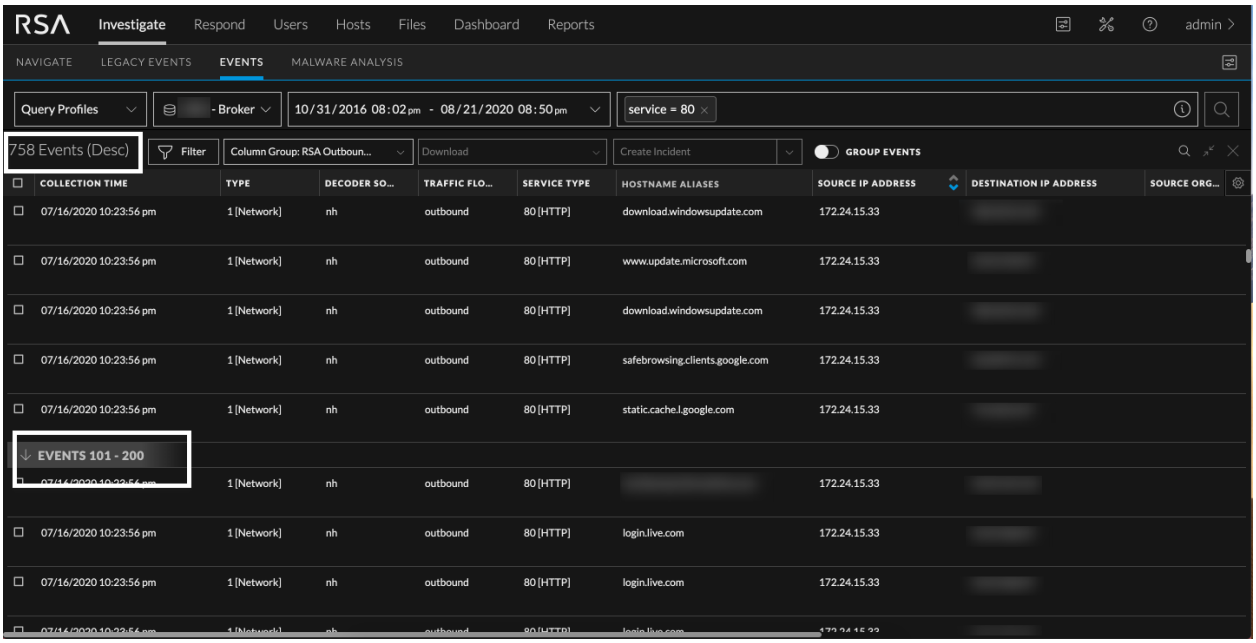
Note: In Version 11.4, the Event Analysis view was renamed as the Events view, replacing the Legacy Events view as the default view analyzing events. Information regarding Events view features prior to version 11.4 also applies to the 11.3 and earlier Event Analysis view.
The Legacy Events view is disabled by default, but the administrator can enable it as described in "Configure Investigation Settings" in the *System Configuration Guide*.

After a query is submitted in the Events view, the Events panel opens with a list of sequential events. The events listed here meet two conditions:

- They match the submitted query.
- They include a value for one or more meta keys required by the selected column group. If you change the column group while viewing the Events list, the original query with the new column group is resubmitted. Unsubmitted query changes made to service, time range, or filter, are ignored.

How Results Are Loaded and Sorted

There is a configurable limit on the number of events that can be loaded; the default value is 5,000. Administrators can configure the limit as described in the *System Configuration Guide*. The events begin loading into the Events panel; a progress bar at the top of the list tracks progress while events are loaded. Events with the earliest collection time are loaded first and a row number indicator of the form " EVENTS xxx - xxx" is inserted in the list after every group of 100 events as shown in the following figure.



A spinner is displayed while the events are loaded. If that count is greater than or equal to the threshold, a message under the spinner advises that fact and directs you to the Query Console for more details. As data begins loading, the message is removed and the spinner remains until all events are loaded. When all events have loaded, one of these messages is added to the bottom of the list:

- "All events loaded."
- "Reached the 5,000 event limit. Consider refining your query."
- "Retrieved 4,000 of 5,000 events prior to query cancellation."

A message at the top of the list indicates the total number of events loaded, if the 5,000 event limit has been reached, and the sorting method in effect:

- The message when fewer than 5,000 events are listed is: "xx,xxx Events"
- The message when more than 5,000 events are listed is: "Oldest 10,000 Events (Asc)"

If the number of events that match the query exceeds the limit of 5,000 events, the newest or oldest 5,000 events in the time window are loaded in ascending order. The portion of events loaded is based on the sort order. For example, if 300,000 events match your query and the sort sequence is set to Ascending, the oldest 5,000 events are loaded by default. You can change this by changing the sort order to Descending and the newest 5,000 events are loaded. Ascending sort, which loads the oldest events first is usually the best setting for investigating network events. If you want to view the newest 5,000 events in the time window, you can change the Default Event Sort Order to Descending in the Event Preferences dialog.

The sorting method for the list is configured in the Event Preferences dialog (see Configure the Events View). Any change in the setting goes into effect the next time you submit a query. The Default Event Sort Order from the Events Preferences dialog is saved in the database and persists after logging out and logging back in.

- **Unsorted** (default for Version 11.4.1): To list events as processed by the Core services. Unsorted is faster because it streams back the events as soon as a match is found versus waiting for all Core services to respond and then displaying them in the chosen order.
- **Ascending** (default for Version 11.4 and earlier): To list the events with the earliest collection time first. The earliest collection time first is well suited to most investigations. When investigating logs, you may want to change the sort sequence to latest collection time first.
- **Descending**: To list the events with the latest collection time first. The latest collection time first is often useful for investigating logs.

Actions to Refine the Events List

When the results are loaded in the Events panel, you can take the actions to refine the list:

- Select a column for sorting events (Use Columns and Column Groups in the Events List).
- Select a set of meta keys (column group) that help with a specific type of investigation (Use Columns and Column Groups in the Events List).
- Apply a query profile (Use Query Profiles to Encapsulate Common Areas for Investigation).
- (Version 11.5) Filter events by pivoting through the metadata Drill into Metadata in the Events View (BETA).

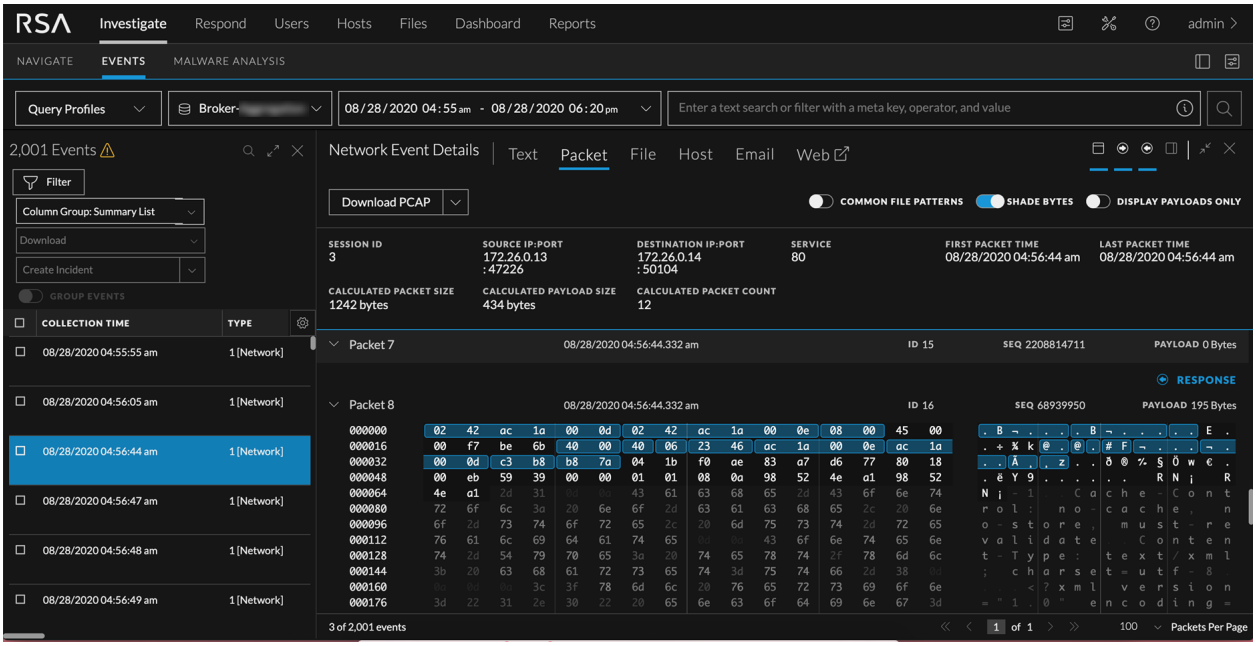
Actions to Analyze Events

The rest of this section has procedures for working in the Events view and adjust reconstructions to bring interesting data into focus.

- You can download events and create an incident in Respond.
- Clicking an event in the Events panel opens the Event Details panel to a tab that shows an event reconstruction (text, packet, file, email, and web) or the tab that shows host information for network events that are enriched with endpoint data (Version 11.5).
- The Events panel and the Event Details panel can be open at the same time.
- Within the Packet tab and the Text tab, you can use additional features to adjust the way the reconstruction is displayed and bring interesting data into focus.

Open, Close, and Adjust the Size of the Panels in the Events View

Initially, the Network Event Details, Log Event Details, or Endpoint Event Details panel occupies 75% of the window width by default.

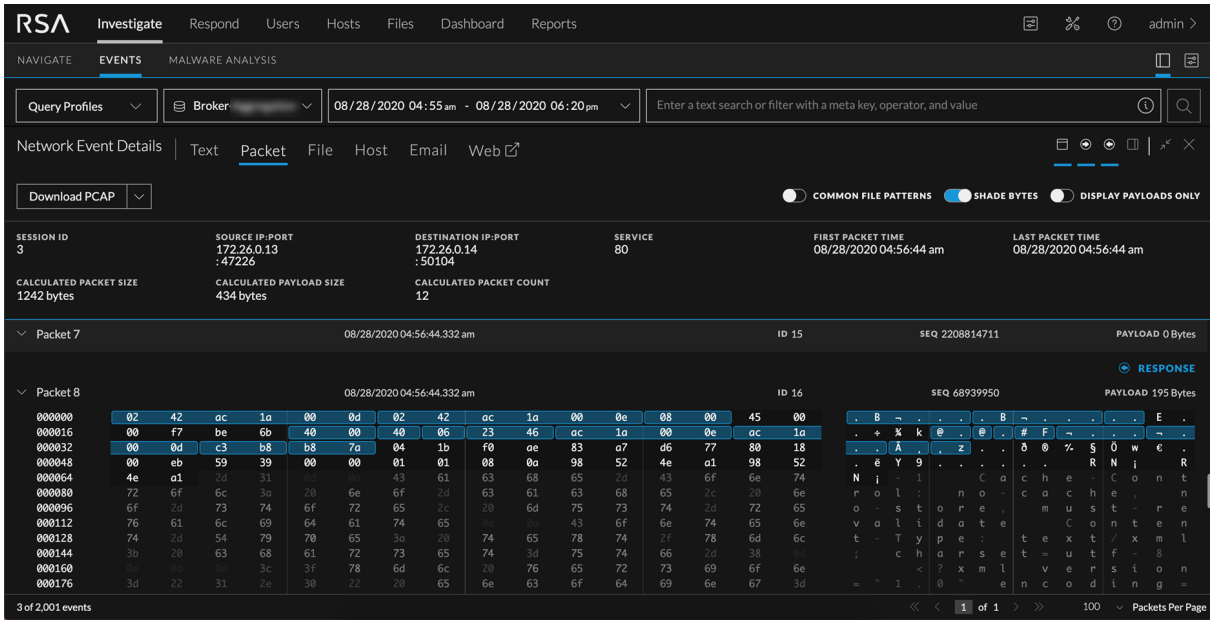


You can adjust the size ratio of the Events panel in relation to the details panels to improve readability by expanding one of the panels, contracting one of the panels, and closing one of the panels. After closing either panel you can reopen it. The ratio you select persists until you change it or refresh the browser.

To optimize your view:

1. To adjust the size ratio of the two panels, do any of the following:
   a. Click ⤢ in the tool bar of the panel that you want to expand.
   b. Click ⤡ in the tool bar of the panel that you want to contract.

2. To close either panel, restoring the open panel to its full width, click .
   This is an example of the reconstruction displayed using the full width of the browser window.



3. To reopen the Events panel after closing, click in the top right corner of the Events view.
   The Events panel opens to the last state (25%-75% or 50%-50%).
4. To reopen the Event Details panel, click an event in the Events panel.

Select the Analysis Type for an Event

To select the analysis type for an event, with an event open in the Event Details panel, click one of the tabs: **Text**, **File**, **Host** , **Packet**, **Email**, or **Web**.

- If you chose **Hosts**, the host information from extended endpoint data is displayed.
- If you chose **File**, **Text**, **Packet**, or **Email**, the reconstruction is displayed.
- If you chose **Web**, the reconstruction of the single event opens in a new tab. This is the same reconstruction of a session used in the Legacy Events view (see Reconstruct an Event in the Legacy Events View).

Note: The packet reconstruction is only available for network events.

Adjust the Display of Requests and Responses

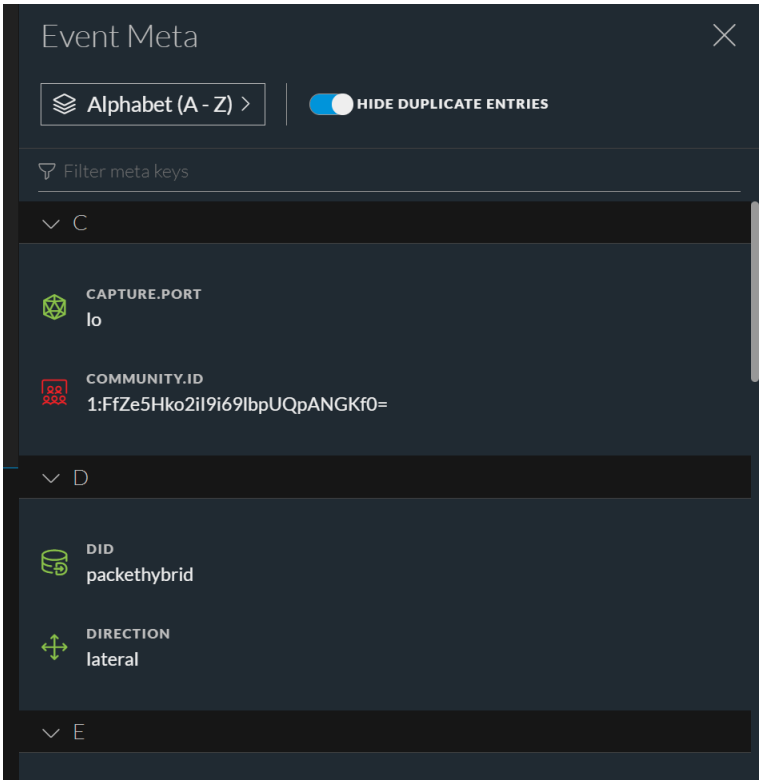
For analysis types that have requests and responses, you can select which side of the conversation to show, Request , Response , or both, click one or both of the direction icons. The reconstruction is refreshed with the selected information.

Note: If you do not see any data, you may have deselected both Request and Response. You must select one of the two to see data displayed.

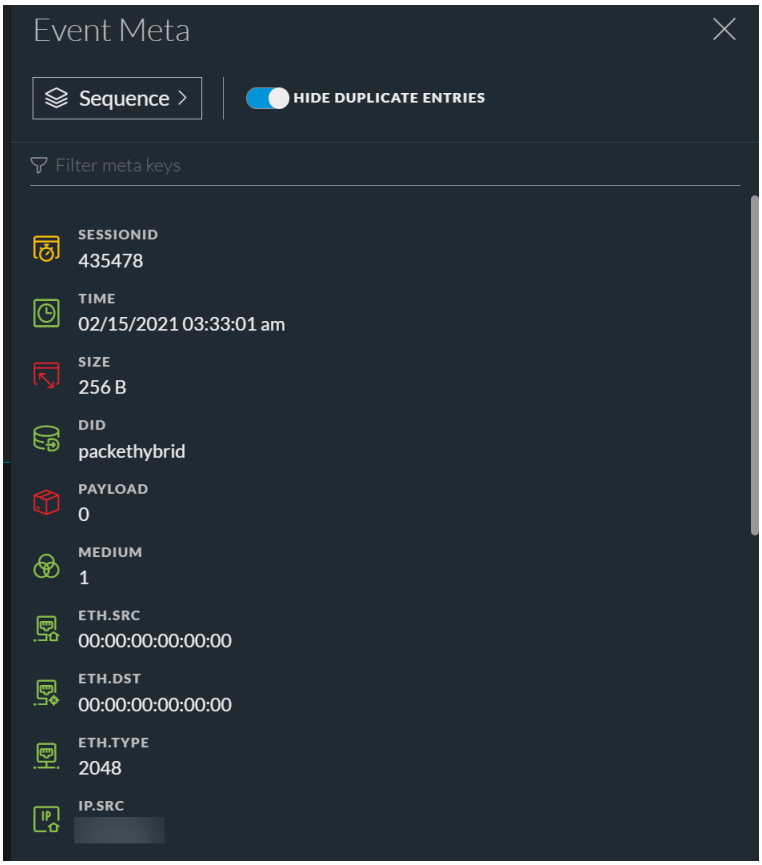
View Associated Metadata for an Event

When examining events in the Text tab, Packet tab, or Filetab, you can click  to show the associated metadata in an adjacent panel, the Event Meta panel.You can change the order of the metadata listed in the Event Meta panel to better find what you are looking for. Metadata can be organized by the sequence in which they were generated or alphabetically by meta key. This figure illustrates metadata organized as meta keys sorted alphabetically."

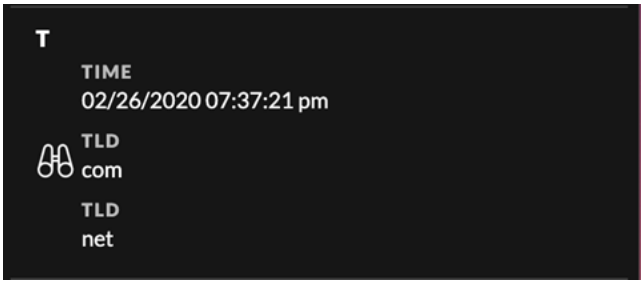
You can change the order of the metadata listed in the Event Meta panel to better find what you are looking for. Metadata can be organized by the sequence in which they were generated or alphabetically by meta key. This figure illustrates metadata organized as meta keys sorted alphabetically.



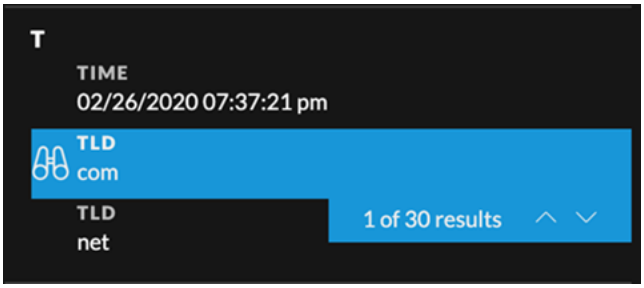
This figure illustrates the same metadata presented in the sequence in which it was generated.



When viewing a text reconstruction and the Event Meta panel, hovering over the meta key/meta value pairs reveals a pair of binoculars if the meta value is searchable in the raw text. This figure is an example of white binoculars on a black background marking a searchable meta key.



Clicking the icon triggers a search for the meta key/meta value pair (case-insensitive) in the Text tab and each instance is highlighted. This is an example of white binoculars on a blue background after clicking a searchable meta key/ meta value combination.

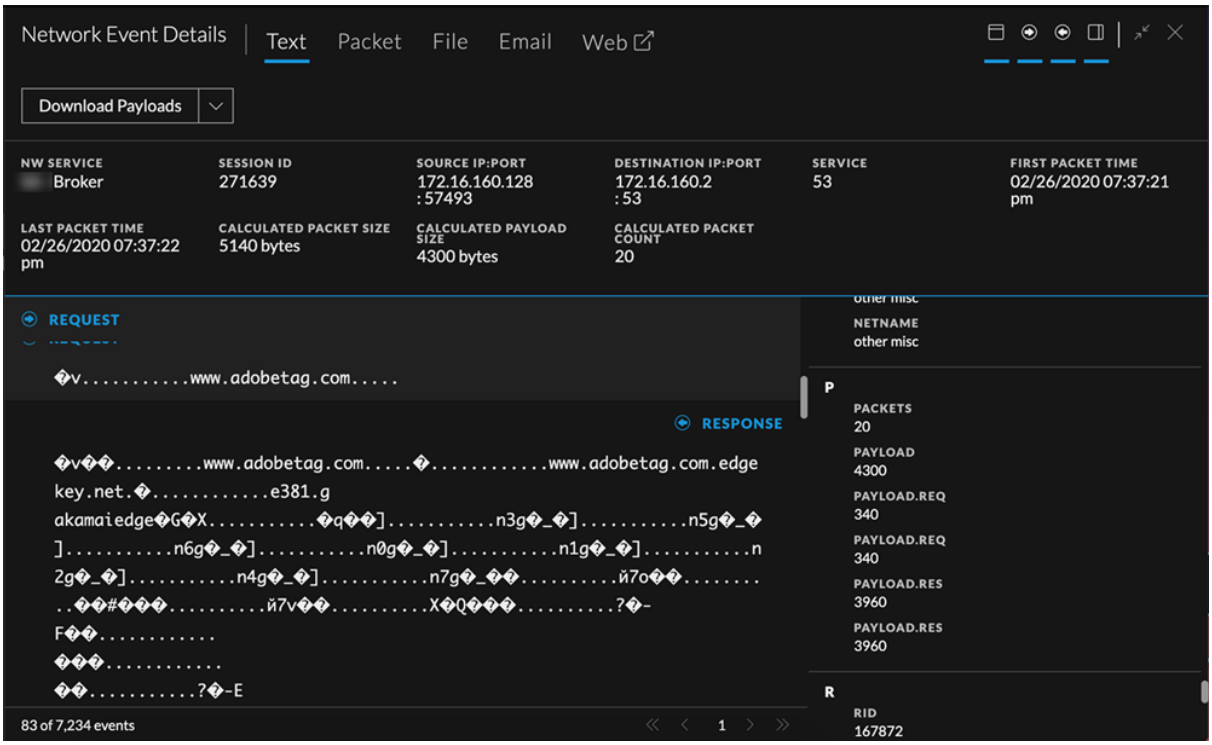


In the Event Meta panel, the highlighted row has a count of the results and up and down arrows that you can use to quickly find each result in the Text tab. You can view each highlighted location of the data that triggered generation of the meta key, going forward to view the next, and back to view the previous.

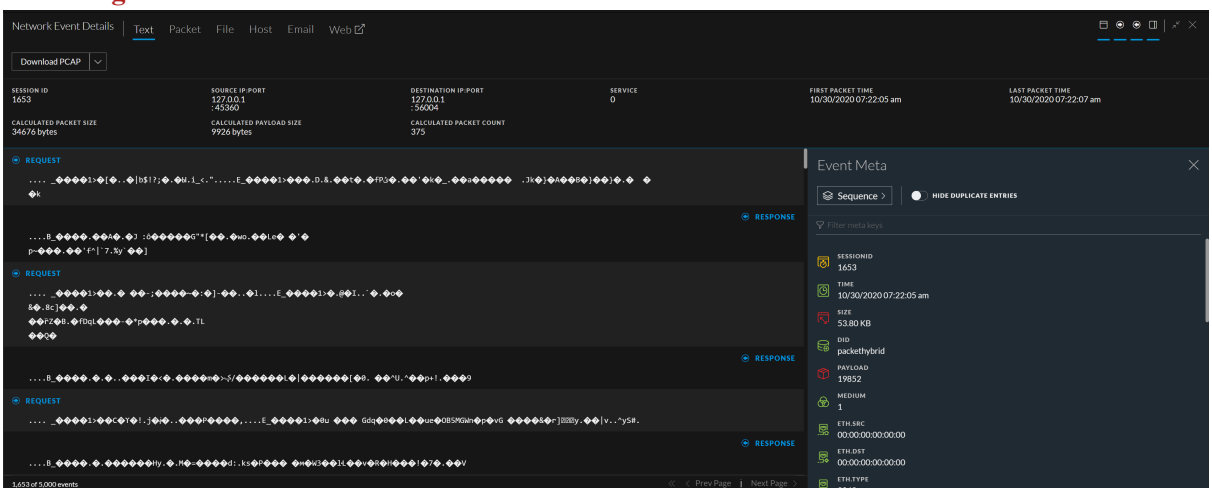
Only meta keys that have relevant values inside the RAW text are searchable. You can search only one meta key at a time. If the value is currently hidden due to truncation of a text entry with more than 3000 characters, the text entry is expanded to reveal the found meta value.

To search the raw text for meta values that triggered a meta key:

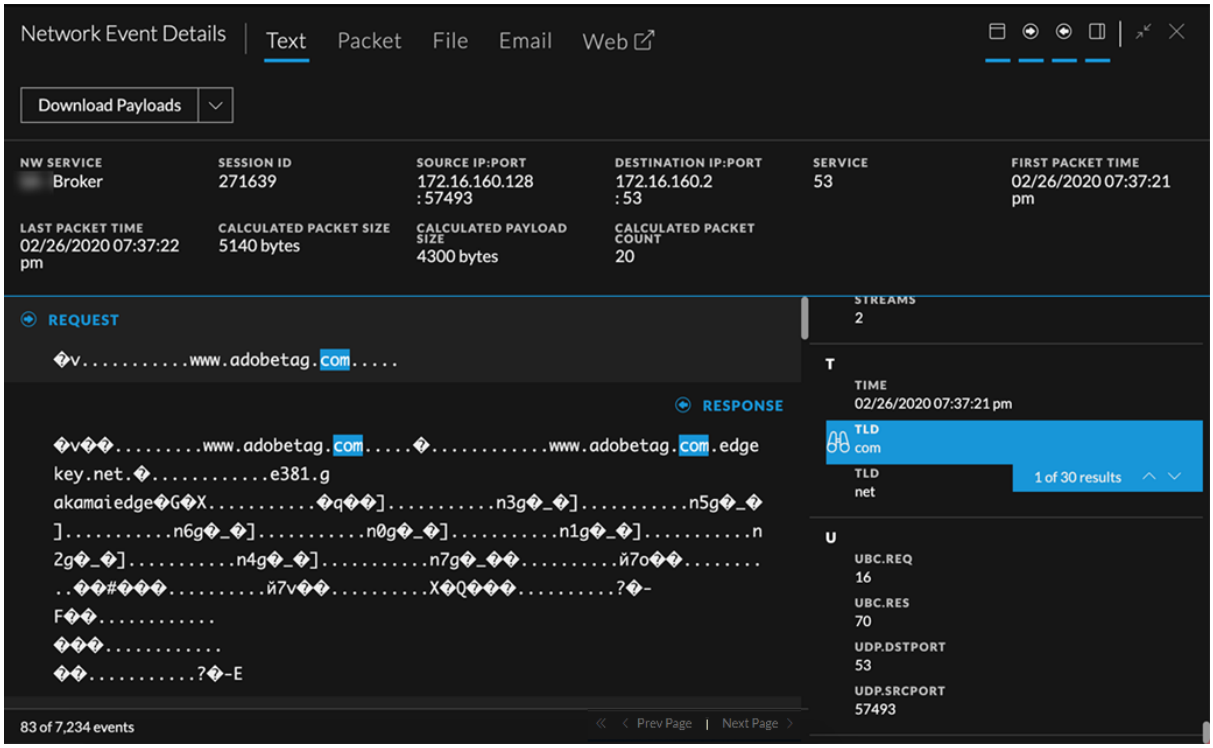
1. Open a network event in the Text tab and click  to open the Event Meta panel.



11.6 Image



2. Hover the mouse over the meta key/meta value pairs in the list until you see a binoculars icon next to a meta key.
3. To search for the value in the raw text, click a row that has the binoculars icon, indicating it is searchable.
   If no relevant occurrence of the value is in the text, the value that you are searching for is highlighted in the Event Meta panel and nothing is highlighted in the Text tab.
   If one or more relevant instances of the value is found in the Text tab, each occurrence is highlighted. The value that you are searching for is highlighted in the Event Meta panel and the up/down arrows for scrolling are visible.



4. To remove the highlighting, click the same meta key/meta value pair in the Event Meta panel, click a different meta key/meta value pair in the Event Meta panel, or close the Event Meta panel.
   The highlighting is removed from the raw text.

Note: When a meta value is more than 255 characters, you can hover over that meta key to view the complete value.

Show or Hide the Event Header

To hide the Event Header in the Packet tab, Text tab, or File tab, providing more vertical space for the data, click  . Clicking the icon again shows the Event Header.

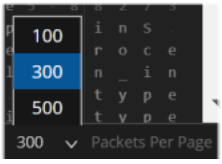
Page Through Events in the Packet and Text Tabs

Pagination controls allow more flexibility in paging through a list of packets or text. In the Packet tab, you can select the number of packets to display per page, and your selection is saved across logins to the NetWitness Platform application. When a control is unavailable, the control is dimmed; for example, when you are viewing page 1, the ◀ and ⏪ controls are dimmed.

Note: Pagination controls are available in Version 11.2 and later of the Text tab. When in the Text tab, you must navigate manually to the last page before the last page control icon is available.

To use pagination controls:

1. With an event open in the Events view, click the current number of packets per page (**50, 100**, **300**, or **500**), and select the new number of packets per page from the drop-down menu.



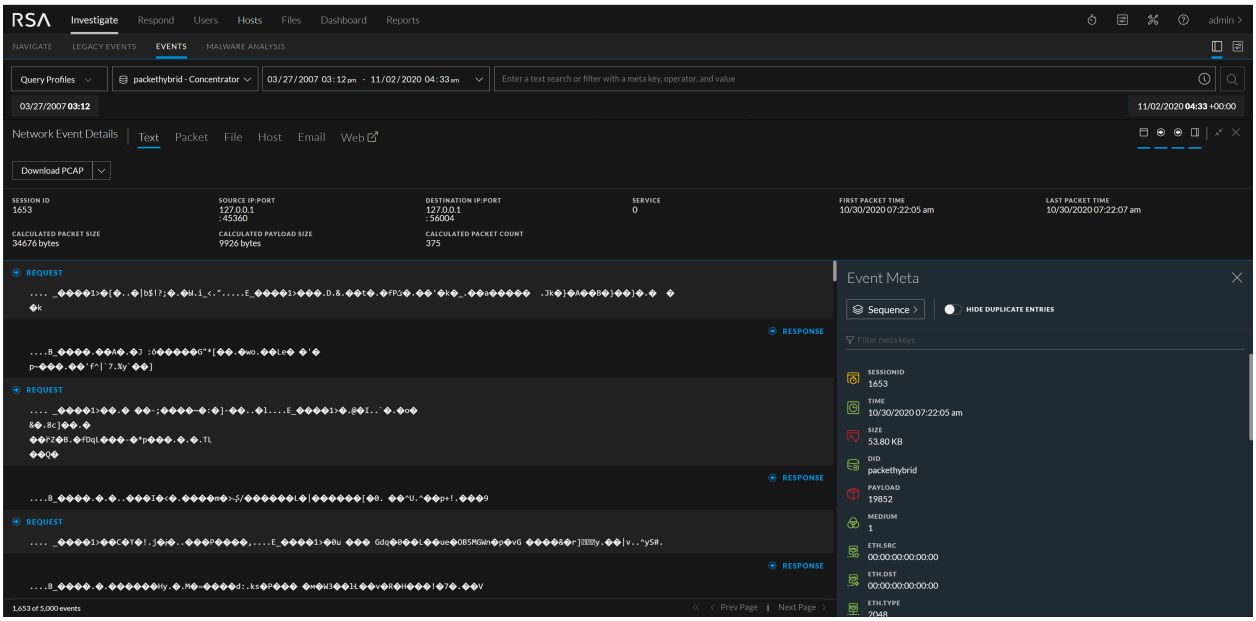
2. To page forward or back, use the page control icons:
   Click ▶ to go to the next page.
   Click ⏩ to go to the last page.
   Click ◀ to go the previous page.
   Click ⏪ to go to the first page.

3. To go to a specific page, type a page number in the page number field 

.

Expand Truncated Text Entries in the Text Tab

A reconstruction of a network event in the Text tab may include requests and responses of many hundred thousands of characters, and scrolling through a long entry that is not of interest can waste time. To save time, text entries that have more than 6000 characters are truncated to show only the first 2000 characters. This example shows an entry that has more than 2000 characters, and a message in the header indicates the percentage of total characters that is being displayed.



You can see that 46% of the characters (the first 2000) are displayed, and click **Show Remaining 54%** to reveal the rest of the entry.

If you search for metadata seen in the Event Meta panel while text is truncated in the Text tab, the truncated text is searched. If the metadata exists inside hidden text, the text entry expands to reveal the text with the found metadata.

Perform URL and Base64 Encoding and Decoding in the Text Tab

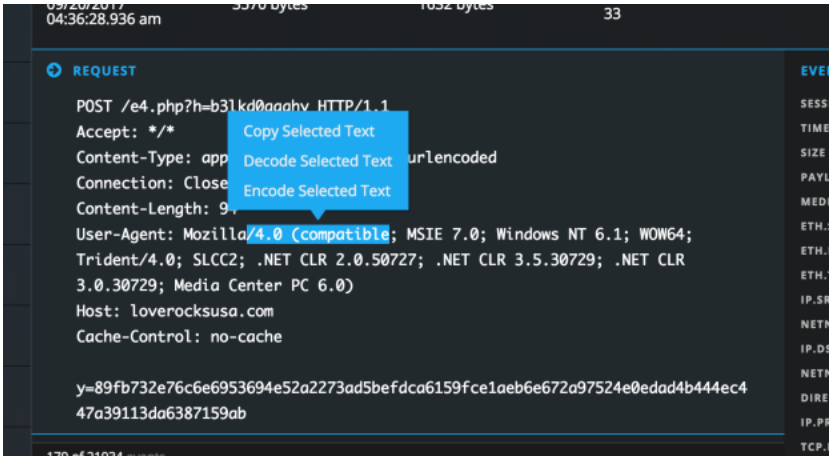
If a network session being reconstructed in the Text tab contains Base64 or URL encoded strings, you can decode a string to better understand the session. If the session contains decoded strings for Base64 or URL, you can view a string in its encoded form in order to search for additional instances of the encoded text in other sessions.

When viewing any network session that contains encoded text in the Text tab, you can select a subset of the text within a single Request or Response to view in either encoded or decoded form. Depending on the content loaded on the Decoder, there may be additional metadata outlining that Base64 or URL encoded data is contained within the session.

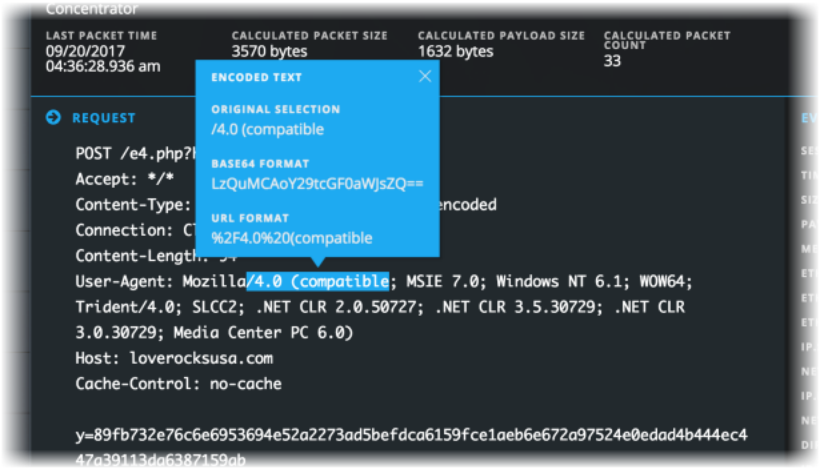
To perform encoding and decoding in the Text tab:

1. In the **Events view**, go to the Text reconstruction of a session that contains encoded or decoded content.

2. To view some decoded text in encoded form, drag to select the text within a Request or Response.
   A menu offers options to encode and decode.



3. Click **Encode Selected Text**.
   The encoded text is displayed in a hover box, which remains in place until you click ☒, select different text in the Text tab, close the Events panel, select another event for reconstruction, or switch to a different reconstruction view.

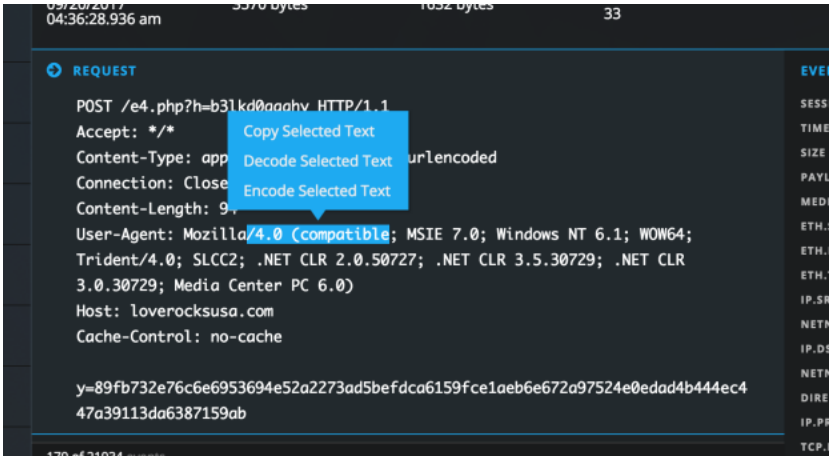


   When a longer text is selected, the hover box is scrollable and large enough to fit the entire selected text and the decoded text.

4. If the session contains encoded text that you want to see in decoded form, drag to select the text within a Request or Response.
   A menu offers options to encode and decode.

5. Click **Decode Selected Text.**
   The decoded text is displayed in a hover box, which remains in place until you click ☒, select different text in the Text tab, close the Events panel, select another event for reconstruction, or switch to a different tab in the Event Details panel.

6. If you want to copy some text from the text reconstruction, do one of the following:
   a. Drag to select some text, right-click, and select **Copy Selected Text** from the pop-up menu.



   b. Drag to select some text, then select either **Decode Selected Text** or **Encode Selected Text**.
   Select the desired text and type **Control-C**.
   The selected text is copied to the clipboard and available to paste in a query.

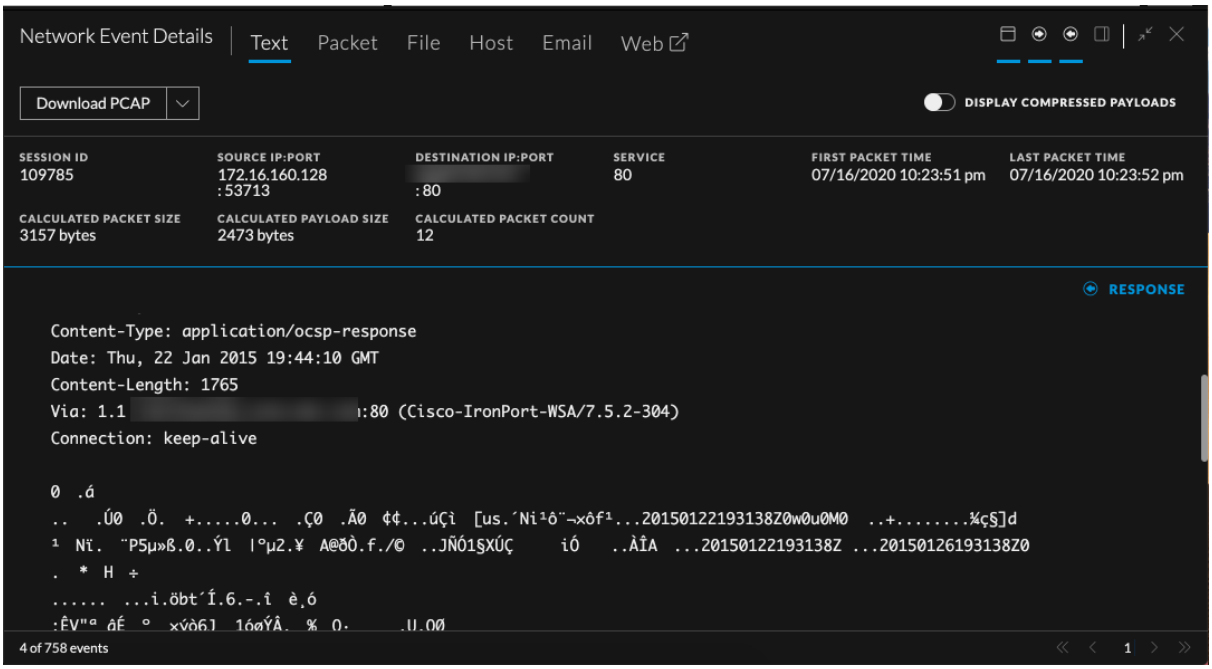
7. When finished, click ☒ to close the hover box.

View Decompressed Text in an HTTP Network Session in the Text Tab

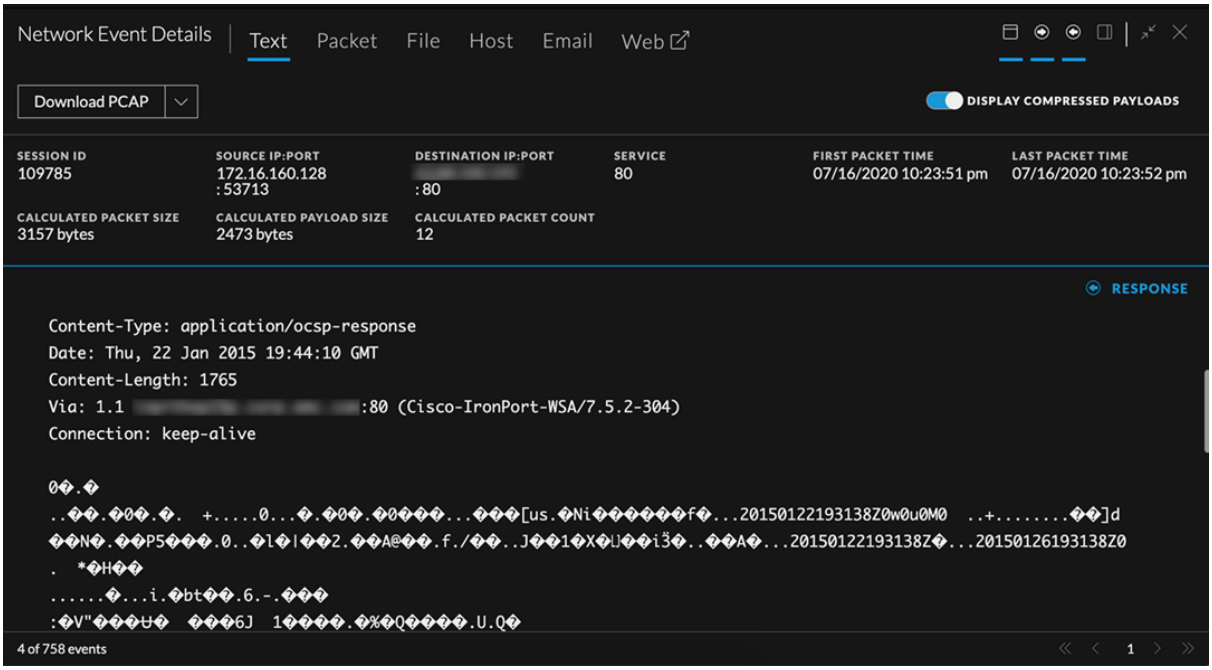
When the content of an HTTP network session is compressed and you are viewing the Text tab, NetWitness displays decompressed content by default. This helps you to determine if there are any patterns and view the readable characters. You can switch between a compressed and decompressed view of compressed text.

The toggle for changing between compressed and decompressed text is only displayed in the Text tab, and is enabled only if there is compressed text content.

1. Open the Text tab of an HTTP session that contains compressed content.
   By default the session is reconstructed with the text decompressed, and above the reconstruction, is the **Display Compressed Payloads** toggle switch.



2. To view the same text in its compressed form, click the toggle switch.
   The view changes so that the compressed text is no longer readable, and the switch indicates the Display Compressed Packets is on.

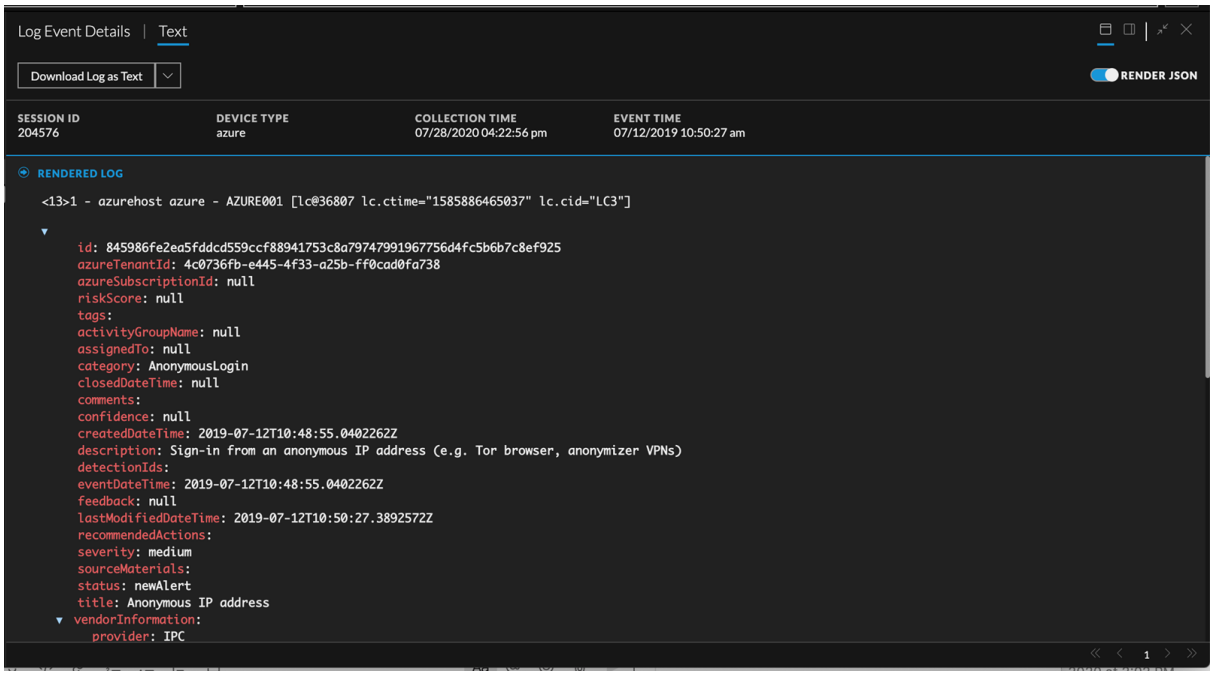


3. To return to the view of decompressed text, click the switch again.

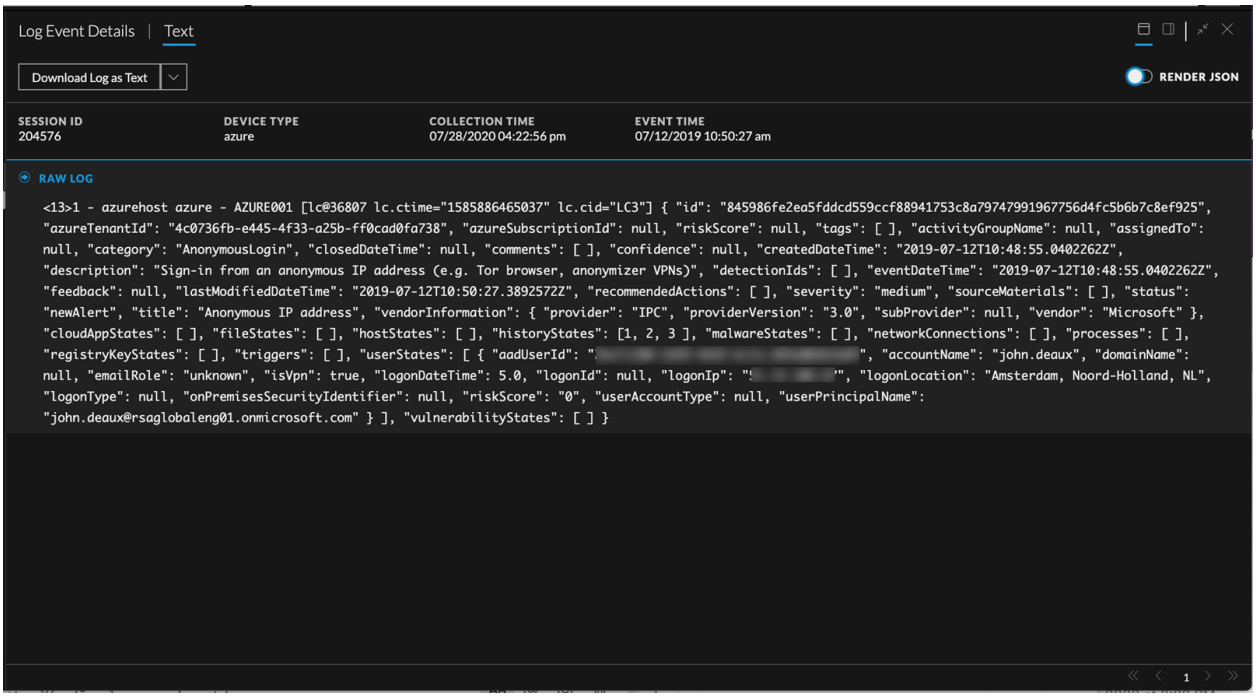
View a JSON String in Tree Format in the Text Tab (Version 11.5.1)

You can view a text reconstruction of a log event in easy-to-read JSON format instead of the raw block format using the Render JSON toggle switch. By default, the switch is enabled, and JSON snippets in a log event are detected and displayed in fully expanded tree format. Invalid JSON snippets are presented as raw text. If you change the setting of the switch, your setting persists in local storage.

1. In the **Events** view, open a log event in Text tab.
   If the raw log contains JSON strings and the Render JSON switch is enabled, all JSON strings found are rendered in tree format.



2. If you prefer to see the log as raw text, click the Render JSON switch.
   The log is rendered as a single block of text without the nested indentation. The setting persists until you change it and will be in the same state when you log in again.

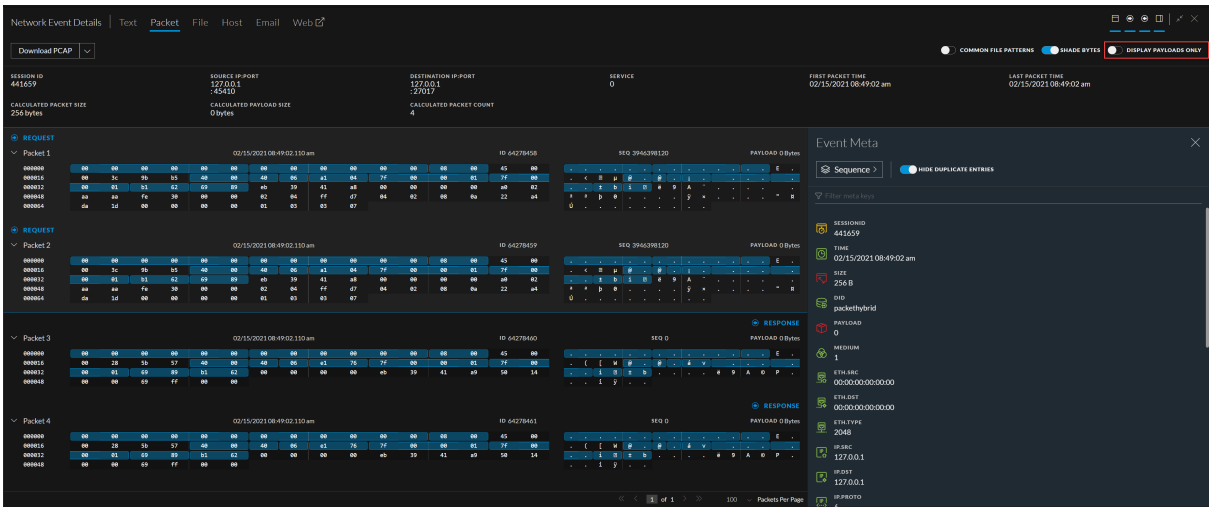


Use the Payload Only Option in the Packet Tab

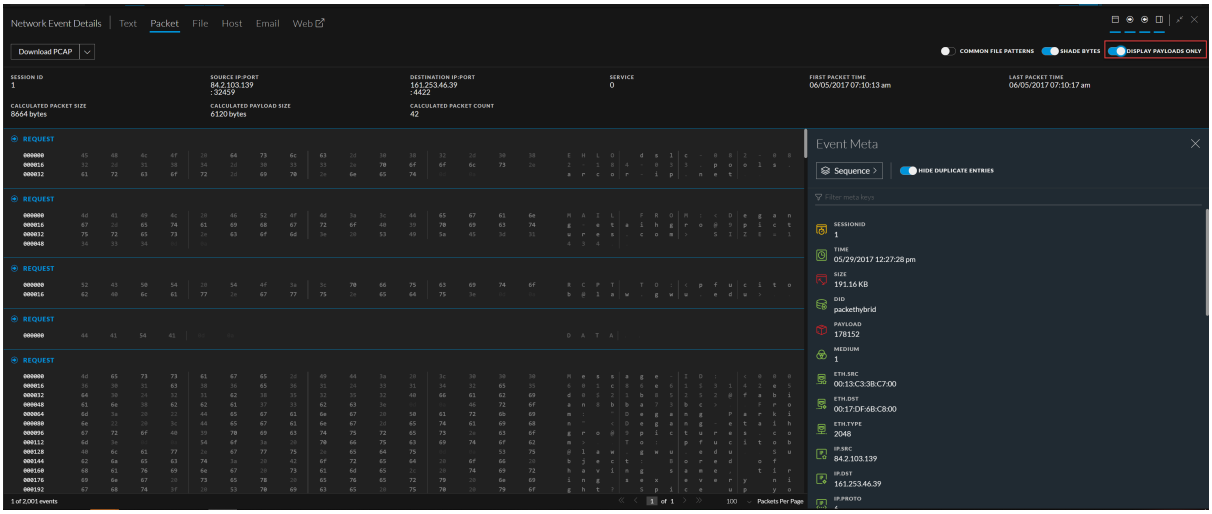
When viewing a reconstruction of a network session in the Packet panel, you can hide the header and footer bytes. The view changes so that only the payload is visible and contiguous same-side packets are concatenated together to make the payload more readable and understandable. This setting persists until you change it or refresh the browser.

- With the Display Payloads Only option off, the number of packets, packet header, packet footer, and payload are displayed.
- With the Display Payloads Only option on, no packet header and footer bytes are displayed. Only the packet content of 16 hexadecimal bytes per line and the corresponding ASCII per line is displayed.

1. In the **Events** view, go to the Packet tab of a network session.
   By default the session is reconstructed with the packet header, footer, and payload displayed.

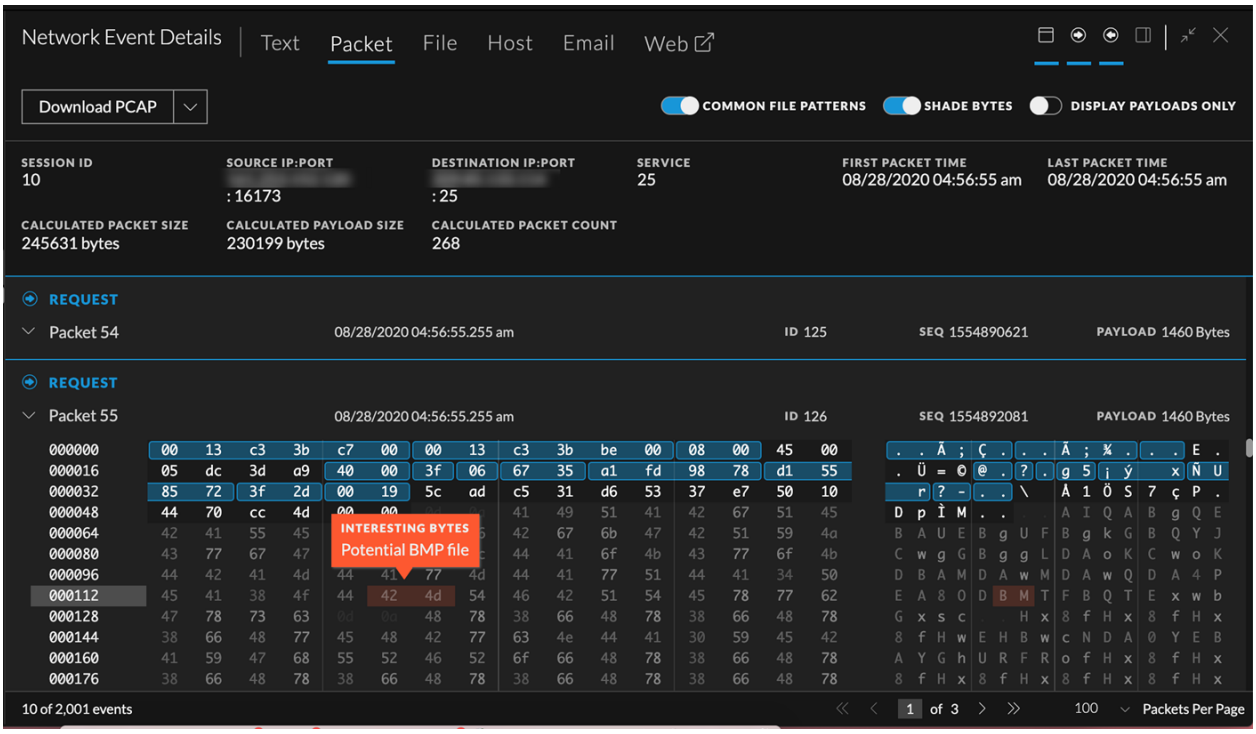


2. To change the view to show only the payload for each packet, click the **Display Payloads Only** toggle switch.
   The view changes so that only the payload is visible.

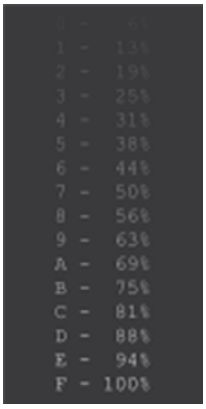


View Highlighted Bytes in the Packet Tab

When you open a reconstruction in the Packet tab, the significant header bytes in each packet are highlighted in blue, and the payload bytes are distinguished using shading to help you understand the contents of the packet. This figure shows the default appearance of the packet reconstruction with highlighting and byte shading.



The Shade Bytes option adds shading to identify the different hexadecimal bytes (00 to FF) using degrees of highlighting. Bytes near the lower range are more transparent, and bytes near 255 are more opaque. Both hexadecimal and ASCII bytes are shaded. This is an example of the shading applied to each hexadecimal byte.



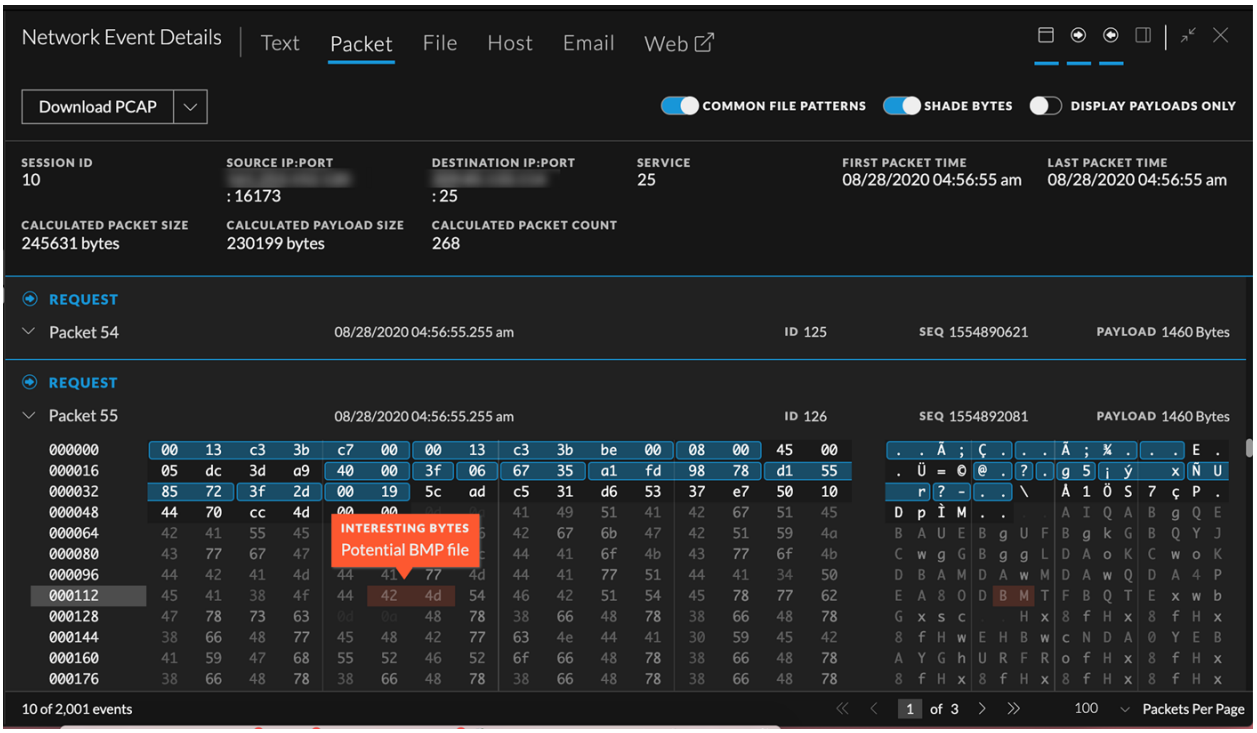
The Shade Bytes switch controls the shading of bytes. When you set Shade Bytes on or off, your setting persists until you change it or refresh the browser.

Highlight Common File Types in the Packet Tab

In the Packet tab, analysts can show or hide highlighting of certain common file types based on the file signature. When the Common File Patterns feature is turned on, the magic number bytes in the file signature are highlighted in the payload and you can hover over the highlighting to see the potential type of file. In this example, `42 4d` is highlighted in the hexadecimal payload and `BM` is highlighted in the ASCII payload. When you hover over the highlighted bytes, the potential file type associated with the magic number is provided in a hover box.

To view common file signatures in the Packet tab:

1. With a reconstruction open in the Packet tab, turn on the **Common File Patterns** option. If there is more then one highlight in view, all are shown.
2. To view the hover box, place the cursor over the highlighting.



These are the file types and corresponding magic numbers that are highlighted if present in the payload:

File Type	Hexadecimal Signature	ASCII Encoding
DOS Executable / Windows PE	4D 5A	MZ
Portable Network Graphics (PNG)	89 50 4E 47 0D 0A 1A 0A	PNG
JPEG	FF D8 FF	JPEG
JPEG/JFIF	4A 46 49 46	JFIF
JPEG/Exif	45 78 69 66	Exif
GIF	47 49 46 38 37 61	GIF87a
GIF	47 49 46 38 39 61	GIF89a
Non-portable Executable	5A 4D	ZM
BMP	42 4D	BM

File Type	Hexadecimal Signature	ASCII Encoding
PDF	25 50 44 46	%PDF
Old Office Document (doc, xls, ppt, msg, and other)	D0 CF 11 E0 A1 B1 1A E1	ÐÏ.à¡±.á
ZIP file formats and formats based on it, such as JAR, ODF, OOXML	50 4B	PK..
7-Zip File Format (7z)	37 7A BC AF 27 1C	7z¼¯'
Java Class File, Mach-O Fat Binary	CA FE BA BE	Êþº¾
Postscript	25 21 50 53	%!PS
Unix/Linux Shell script	23 21	#!
Executable and Linkable Format (ELF) executables	7F 45 4C 46	.ELF

Reconstruct an Event in the Legacy Events View

When viewing a list of events in Legacy Events view, you can safely create a reconstruction of the event in a readable form that matches the original. By default, the initial view of a reconstructed event is the most suitable format (Best Reconstruction); for example, web content is reconstructed as a web page; an IM conversation is displayed with both parts of the conversation. Each user can select a different default reconstruction in the Profile > Preferences view.

You can also open a reconstruction from the Navigate view if you know the Event ID of the event.

In the reconstruction, you can:

- Select event information to view. Possible values are: request data, response data, both request and response data.
- Select the reconstruction type: details, text, hex, packets, web, mail, or IM.
- Export raw logs.
- Export the event as a PCAP file.
- Extract any files available in the event.
- Extract all the meta data associated with the event.

> **Caution:** Be careful when clicking a link to a file in the Reconstruction. If your system has an application associated with the file, or the browser is capable of opening them, and the attachments are malicious, they can negatively affect your system.

- Display the event in a separate window or tab (depending on your browser configuration).
- If you are viewing the reconstruction as a preview in the current view, you can page forward to the next event and back to the previous using the navigation buttons in the bottom left corner.

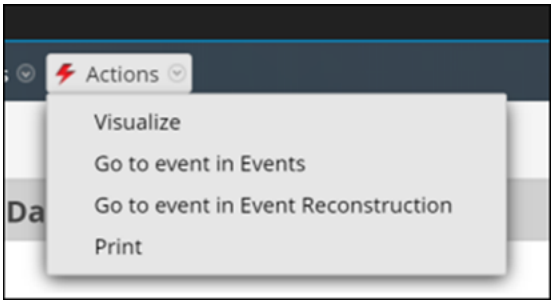
Note: Reconstruction Settings and Reconstruction Cache Settings allow an administrator to manage application performance for Investigation (as described in the *System Configuration Guide*). When analysts reconstruct sessions, two situations can affect performance and results.
- Some large events contain many thousands of source packets. Reconstructing these sessions can degrade application performance.
- In some cases, the reconstruction cache can present incorrect content; for this reason, NetWitness cleans cache that is older than a day every 24 hours. Between the daily cache cleanings, certain actions my result in stale cache being used for a reconstruction, and if the need arises, administrators can manually clear cache for one or more services that are connected to the current NetWitness Server.

Reconstruct an Event Using an Event ID

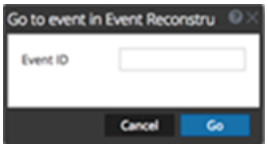
You can reconstruct an event directly from the Navigate view given a known Event ID. You can use this option without executing a query as you usually do when beginning an investigation. A service and time range must be selected to be able to jump directly to an event using just its `eventid`.

To view a reconstruction or event analysis directly from the Navigate view:

1. Go to **Investigate > Navigate** and select **Actions > Go to event in Events** or **Go to event in Event Reconstruction**.



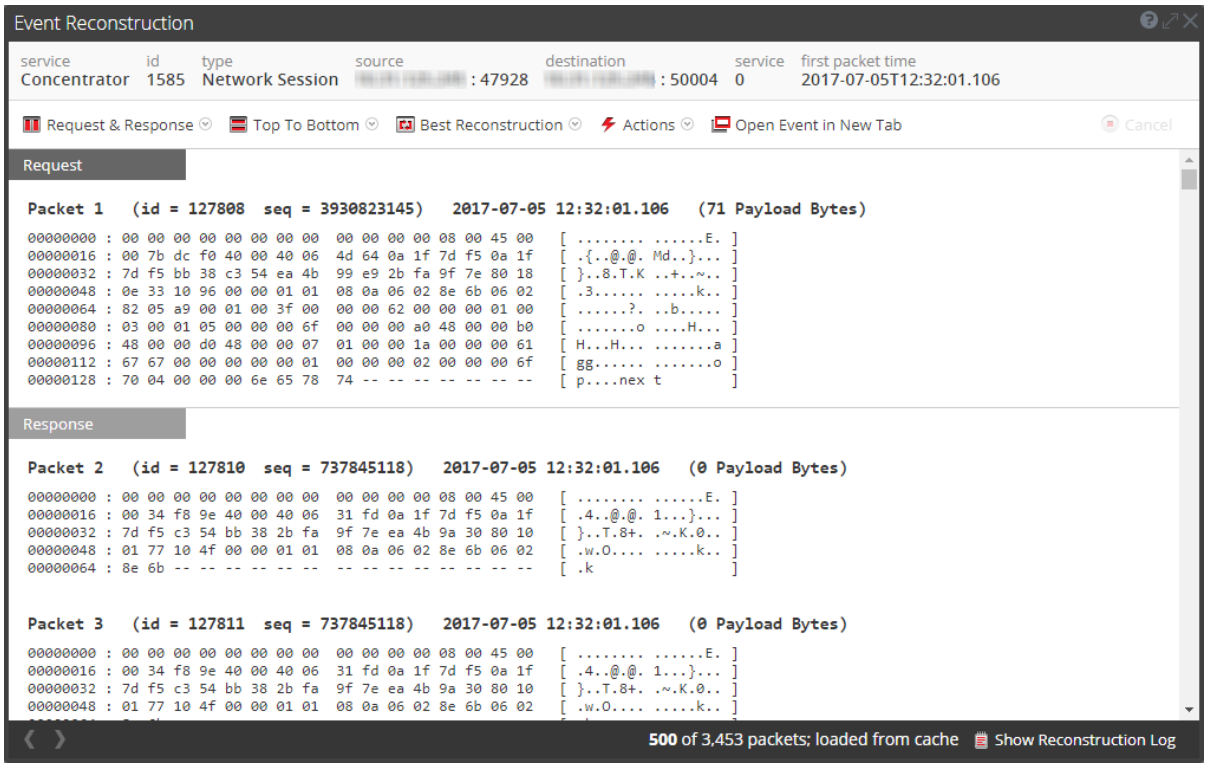
The Go to event dialog is displayed. There are two dialogs, one for Events and one for Legacy Event reconstruction. Both ask for the Event ID.



2. In the **Event ID** field, type the ID and click **Go**.
   The specified event is reconstructed in the legacy Event Reconstruction view or the Events view.

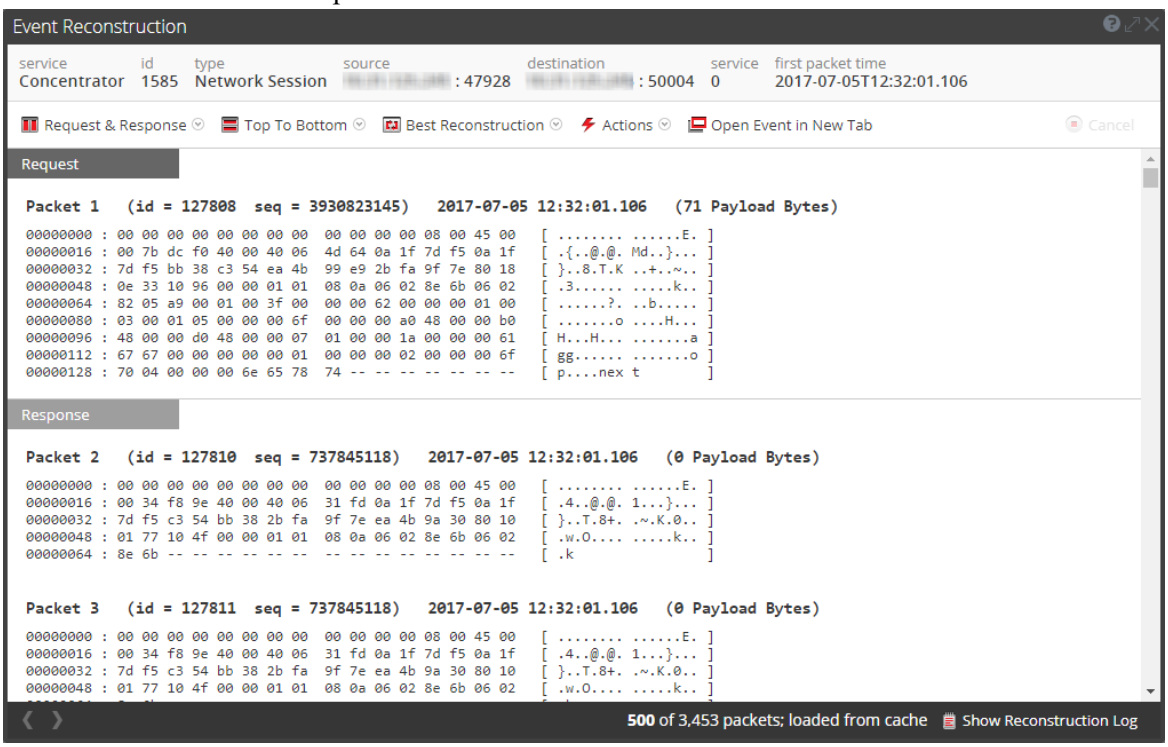
Reconstruct an Event from a Drill Point in the Navigate View

1. Click the count (the green number following a value) for a value in the Navigate view to open a drill point in the **Events** view.
2. To show all meta data, click + Show Additional Meta .
3. To open an event reconstruction in the Legacy Events view, select an event to reconstruct and select **Actions > View Event > Preview Inline**.
   The Event Reconstruction opens in a popup window in the same view. By default, NetWitness displays the best reconstruction for the event determined by the event content or the reconstruction that you have selected in the Default Session View setting for Investigation. You can use the options in the Event Reconstruction toolbar to change the reconstruction method, view side-by-side results, export an event, open an email attachment, extract files, and open the event in a new tab. The toolbar options vary depending on the type of event being reconstructed (network event, log event, or endpoint event). This is an example of the reconstruction for a network event.



4. To preview a reconstruction of the next event, click  in the lower left corner of the reconstruction or to preview a reconstruction of the previous event, click  .
5. To open an event reconstruction in a new tab, do one of the following:
   a. In the **Legacy Events** view, select an event to reconstruct and select **Actions > View Event > Open in New Tab**.
   b. In the **Event Reconstruction** toolbar of previewed reconstruction, click **Open Event in New Tab** in the toolbar.

The Event Reconstruction opens in a new tab.



View Side by Side or Top to Bottom

To select the way requests and responses for an event are displayed:

1. In the **Event Reconstruction** toolbar, click **Top to Bottom** or **Side by Side**.
2. In the drop-down menu, select the information you want to see in the event: **Side by Side** or **Top to Bottom.**
   The reconstruction is refreshed with the selected information.

Select Event Information to View

To select what event information to view:

1. In the **Event Reconstruction** toolbar, click **Request & Response**.
2. In the drop-down menu, select the information you want to see in the event: **Request & Response**, **Request**, or **Response**.
   The reconstruction is refreshed with the selected information.

Select Event Reconstruction Type

To select the reconstruction type for an event:

1. In the **Event Reconstruction** toolbar, click **Best Reconstruction**.
2. In the drop-down menu, select the reconstruction type to view: **meta**, **text**, **hex**, **packets**, **web**, **mail**, or **files**.
   The reconstruction is refreshed with the selected reconstruction type.

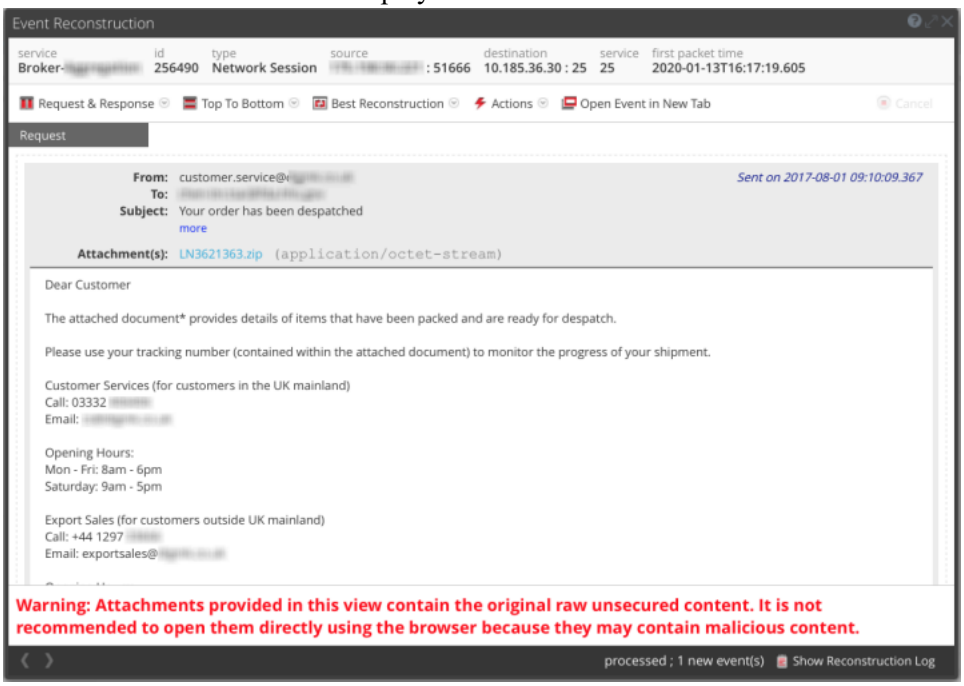
Open or Download an Email Attachment

When viewing a reconstruction of an email that has attachments, you can open supported file types or download the files to the local system.

Caution: Be careful when selecting file attachments. If your system has an application associated with the file attachments, or the browser is capable of opening them, and the attachments are malicious, they can negatively affect your system.

To open or download email attachments:

1. In the **Event Reconstruction** toolbar, select the **View** drop-down and select **View Mail**.
   The Event Reconstruction is displayed.



2. In the **Event Reconstruction** section of the email, click the Attachment.
   If the file type is supported by the browser, the attachment will open in a new tab.
   If the file type is not supported, the Download dialog is displayed so that you can download the attachment.

Export an Event as a PCAP File

The PCAP export option downloads the sessions for the current time range and drill point to a PCAP file. To export an event as a pcap file:

1. In the **Event Reconstruction** toolbar, click **Actions**.
2. Click **Export PCAP**.
3. A confirmation dialog is displayed.
4. Click **OK**.
   The job is scheduled and when complete the PCAP is downloaded to the local file system. In the Profile > Jobs tab, you can download the PCAP.

Extract Files from a Reconstructed Event

The Extract Files option extracts and downloads the files associated with the event. To extract files:

1. In the **Event Reconstruction** toolbar, click **Actions**.
2. Click **Extract Files**.
   The File Extraction dialog is displayed.
3. Select the types of files to extract, and click **OK**.
4. The job is scheduled and when complete the selected file types are downloaded to the local file system. In the Profile > Jobs tab, you can download the files.

Look Up Additional Context for Results

The Context Hub is a centralized service that aggregates data about entities from multiple configurable data sources. This data can extend your investigation with additional context beyond the immediate results of a specific query. For example, the Context Hub can tell you if a given entity has been mentioned in any incidents, alerts, feeds, or community intelligence publications.

To enable viewing of contextual information, your administrator must add the Context Hub service in RSA NetWitness Platform and configure data sources for the Context Hub service as described in the *Context Hub Configuration Guide*. Analysts also need a role with the permission `Context Lookup` as described in "Role Permissions" and "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*.

When the Context Hub service is enabled and configured, NetWitness provides enrichment data from NetWitness Respond, custom lists, and NetWitness Endpoint directly in the Navigate view, Events view, and Legacy Events view. A visual cue highlights meta values for which enrichment data is available in the Investigate views, and you can click on the highlighted value to look up the contextual information and intelligence. You can look up details and intelligence about elements associated with an event in the Context Hub. These elements, or entities, are identifiers, such as an IP address, a user name, a host name, a domain name, a file name, or a file hash. The data from configured sources, such as RSA NetWitness Endpoint, can help you understand what is happening. In Version 11.5 and later, you can add STIX data sources and view related data using context lookup; associated elements are IP address, file name, file hash, domain name, and URL.

In addition, you can add lists and list values for Context Hub enrichment; you can view lists, edit meta values in an existing list, or create a new list. When you add meta values to a list, you can investigate the meta values using the context lookup option.

Note: In Version 11.2 and earlier, you can look up additional context in the Navigate view or the Legacy Events view, but not the Event Analysis view.

For an analyst to manage lists in Investigate, the administrator must:

- Enable the Context Hub service.
- Assign an analyst role with permission `Manage List from Investigation` to the user who will perform Context Lookup from Investigation views.
- Configure appropriate roles and permissions as described in "Role Permissions" and "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*.

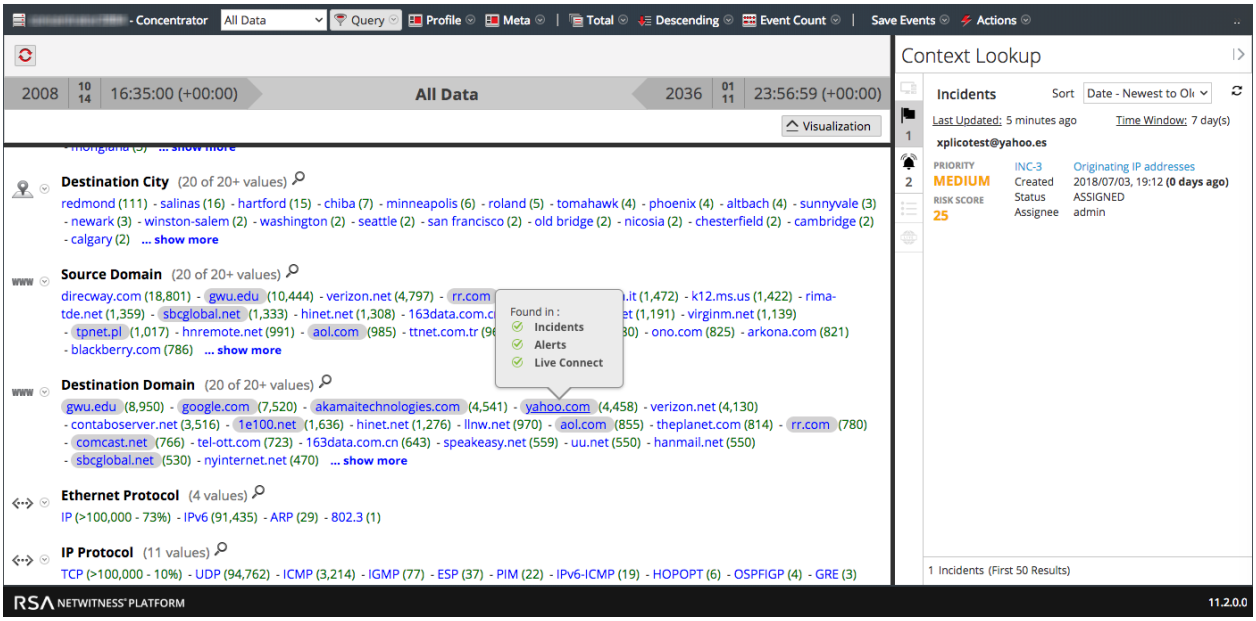
In Version 11.6 and later, you can add REST API data sources and view related data using context lookup. Also, Administrators can now configure specific Context Hub sources (For example, specific lists, Respond, Endpoint, and so on) for context highlighting during investigation. If context highlighting is disabled for a Context Hub source, analysts will see results from all the data sources when opening the Context Panel for a meta value, but the values are not highlighted in the Investigation views. In the View Context:

- No data is shown if the meta values is not highlighted.
- If there are entities common across different data sources, the meta values for those entities are underlined for all the data sources, but the data is shown only for the data source on which context highlighting is enabled.

Open the Context Lookup Panel

In the Context Lookup panel, you can view and explore individual data sources for further investigation. For a detailed description of the information displayed for each data source, see Context Lookup Panel.

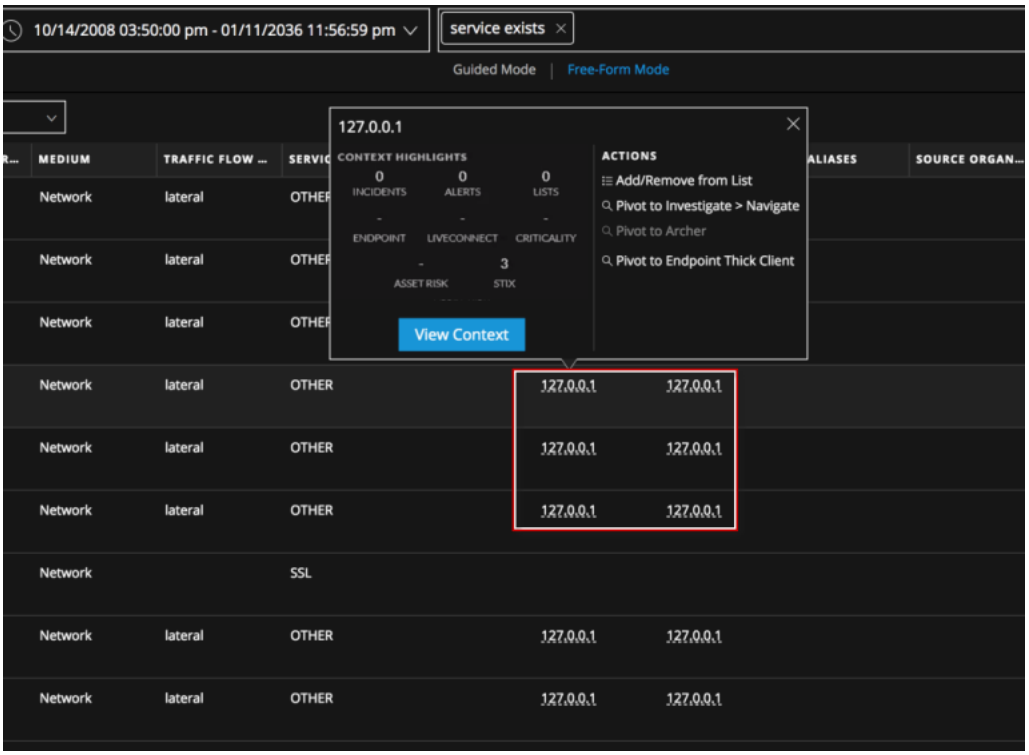
In the Navigate view and Legacy Events view, entities that have associated context data available are highlighted with a gray background; hovering over an entity displays a hover box giving a summary of the available data. When you right-click the entity, the Context Hub queries the configured data sources for relevant information, and the Context Lookup panel opens from the right side of the browser window. The Context Lookup panel is populated with the information from the Context Hub as it becomes available. You can perform another lookup by right-clicking on another entity, and the Context Lookup panel is updated with that entity's information.



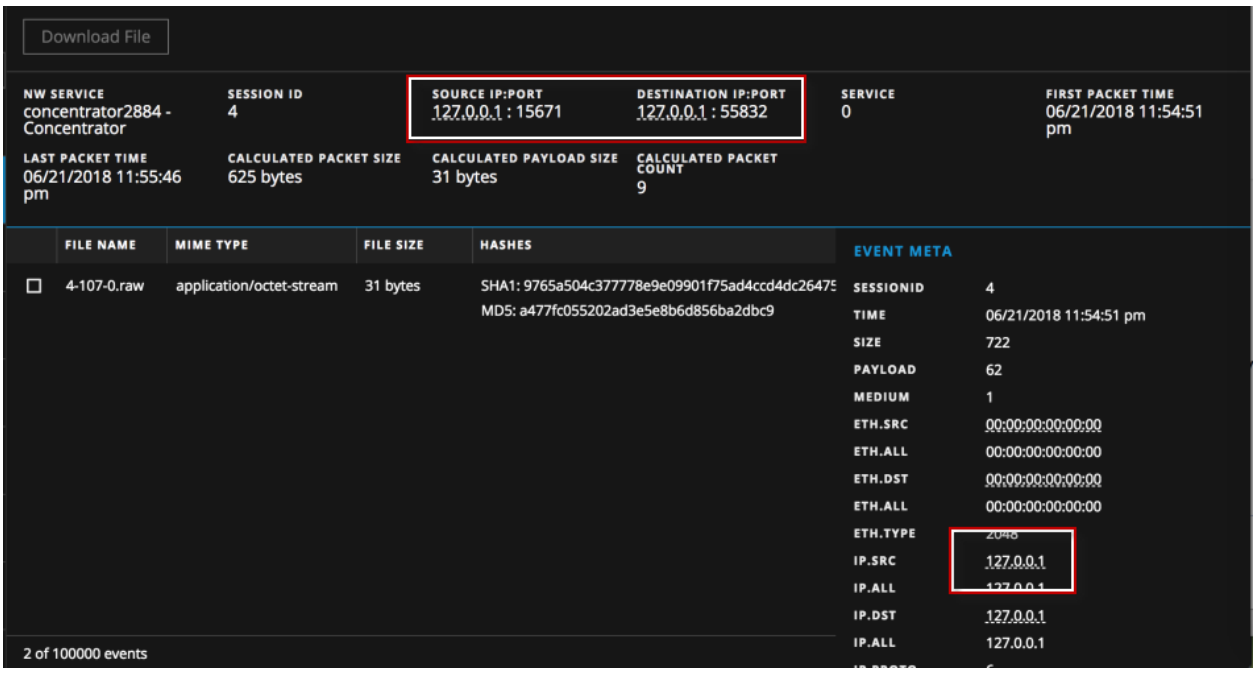
In the Events view, you can see underlined entities in the Events panel, the Event Header, or the Event Meta panel. If an entity is underlined, NetWitness is populating information about that entity type in the Context Hub. There may be additional information available about that entity in the Context Hub.

The following figure shows underlined entities in the Events panel with the context tooltip open. The context tooltip has two sections: Context Highlights and Actions.

- The information in the Context Highlights section helps you to determine the actions that you would like to take. It can show related data for Incidents, Alerts, Lists, Endpoint, Live Connect, Criticality, Asset Risk, and STIX. Depending on your data, you may be able to click these items for more information.
- The **Actions** section lists the available actions. In the example, the Add/Remove from List, Pivot to Investigate > Navigate, Pivot to Archer, and Pivot to Endpoint Thick Client options are available.



The following figure shows underlined entities in the Event Header and the Event Meta panel.

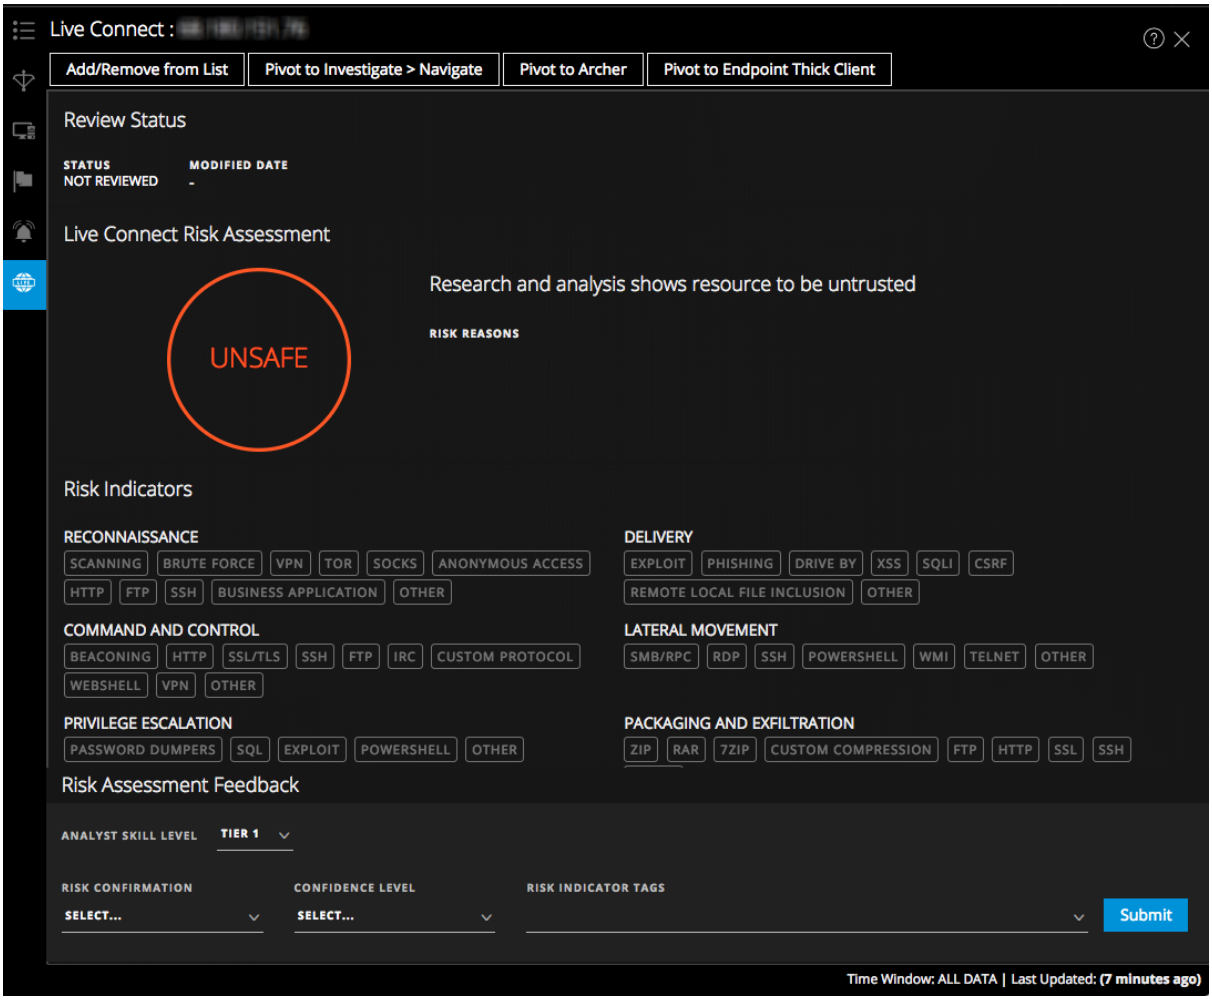


When you click View Context in the context tooltip, the Context Hub queries the configured data sources for relevant information, and the Context Lookup panel opens from the right side of the browser window. The Context Lookup panel is populated with the information from the Context Hub as it becomes available. You can perform another lookup by using the View Context option on another entity, and the Context Lookup panel is updated with that entity's information.

You can also take any available action in the Actions section.

To view information in the Context Lookup panel in the Events view:

1. Hover over different meta values to see the data sources for which data is available.
   A context tooltip displays a list of the context data available for the selected meta value.
2. Click **View Context** in the context tooltip to open the Context Lookup panel.
   The Context Lookup panel opens from the right side of the browser window. The Context Lookup panel is populated with the information from the Context Hub as it becomes available.



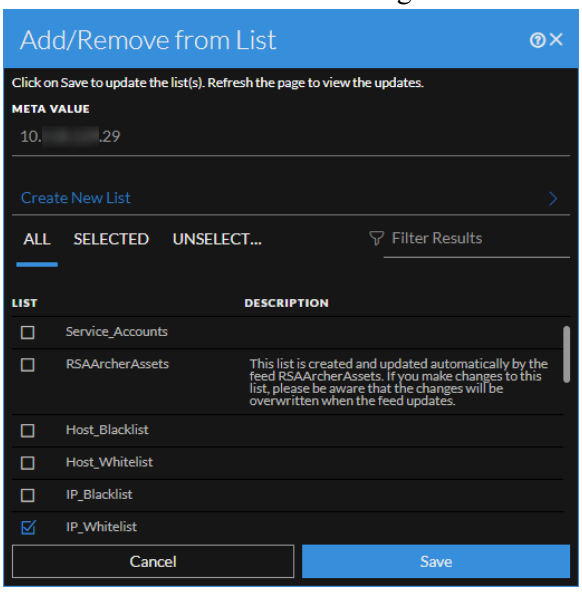
3. To perform actions on an entity, select one of the available actions in the context tooltip: Add /Remove from List, Pivot to Investigate > Navigate, Pivot to Archer, Pivot to Endpoint Thick Client. For more information, see Pivot to Investigate > Navigate (Events View), Pivot to Archer (Events View), Pivot to NetWitness Endpoint Thick Client (Events View), and Add an Entity to a Whitelist.

Note: The Pivot to Archer link is disabled when Archer data is not available or when the Archer data source is not responding. Check that the RSA Archer configuration is enabled and configured properly.

Add an Entity to a Whitelist

You can add any underlined entity to a list, such as a Whitelist or Blacklist, from a context tooltip. For example, to reduce false positives, you may want to whitelist an underlined domain to exclude it from the related entities.

1. In the Events panel, the Event Header, or the Event Meta panel, hover over the underlined entity that you would like to add to a Context Hub list.
   A context tooltip showing the available actions is displayed.
2. In the **ACTIONS** section of the tooltip, click **Add/Remove from List**.
   The Add/Remove from List dialog shows the available lists.



3. Select one or more lists and click **Save**.
   The entity is added to the selected lists.

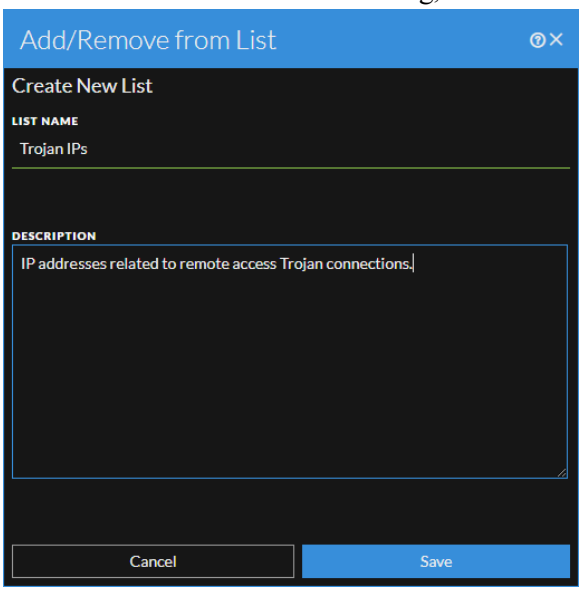
Create a List (Events View)

You can create lists in Context Hub from the Events view. In addition to using lists to whitelist and blacklist entities, you can use lists to monitor entities for abnormal behavior. For example, to improve the visibility of a suspicious IP address and Domain under investigation, you may want to include them in two separate lists. One list could be for domains suspected of being related to command and control connections, and another list could be for IP addresses related to remote access Trojan connections. You can then identify indicators of compromise using these lists.

To create a list in the Context Hub:

1. In the Events panel, the Event Header, or the Event Meta panel, hover over the underlined entity that you would like to add to a Context Hub list.
   A context tooltip showing the available actions is displayed.
2. In the **ACTIONS** section of the tooltip, click **Add/Remove from List**.

3. In the Add/Remove from List dialog, click **Create New List**.



4. Type a unique **List NAME** for the list. The list name is not case sensitive.
5. (Optional) Type a **DESCRIPTION** for the list.
   Analysts with the appropriate permissions can also export lists in CSV format to send to other analysts for further tracking and analysis. The *Context Hub Configuration Guide* provides additional information.

Pivot to Investigate > Navigate (Events View)

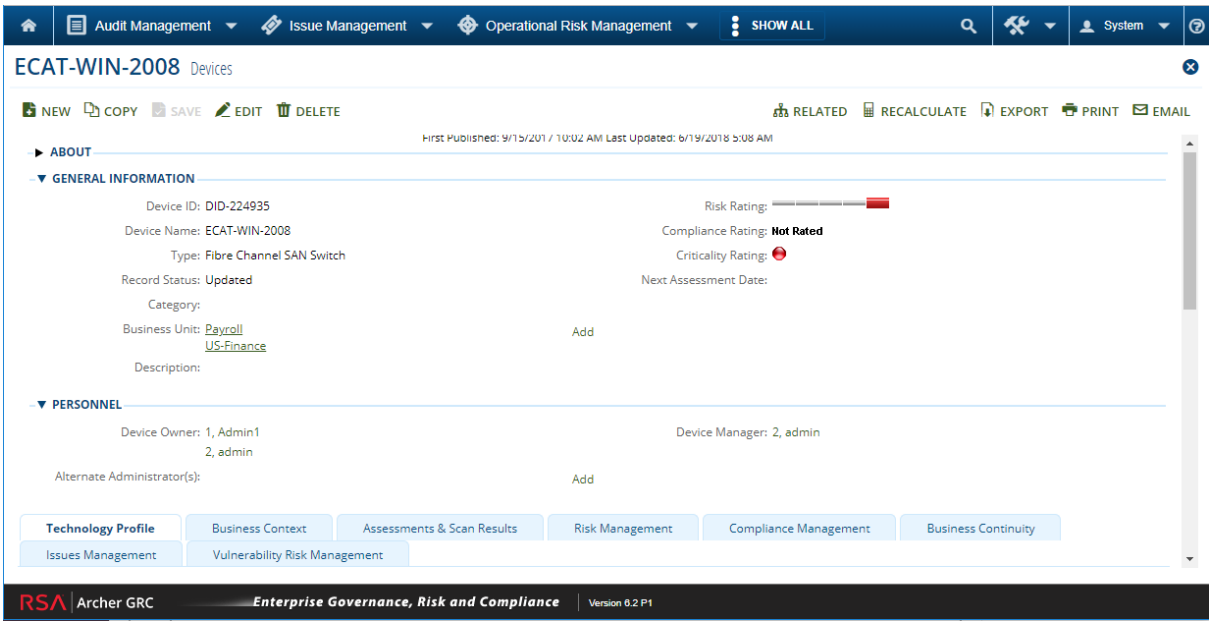
For a more thorough investigation of an entity, you can open the the Navigate view.

1. In the Events panel, the Event Header, or the Event Meta panel, hover over any underlined entity.
2. In the **ACTIONS** section of the tooltip, select **Pivot to Investigate > Navigate**.
   The Navigate view opens, enabling you to perform a deeper dive investigation. For more information, see Begin an Investigation in the Navigate or Legacy Events View.

Pivot to Archer (Events View)

For viewing more details about the device in RSA Archer Cyber Incident & Breach Response, you can pivot to the device details page. This information is displayed only for IP address, host, and Mac address.

1. In the Events panel, the Event Header, or the Event Meta panel, hover over any underlined entity (IP address, host, and Mac address).
2. In the **ACTIONS** section of the context tooltip, select **Pivot to Archer**.
3. The device details page in **RSA Archer Cyber Incident & Breach Response** opens if you are logged in to the application, otherwise the login screen is displayed.



Note: The Pivot to Archer link is disabled when Archer data is not available or when the Archer Datasource is not responding. Check that the RSA Archer configuration is enabled and configured properly.

For more information, see the *Archer Integration Guide*.

Pivot to NetWitness Endpoint Thick Client (Events View)

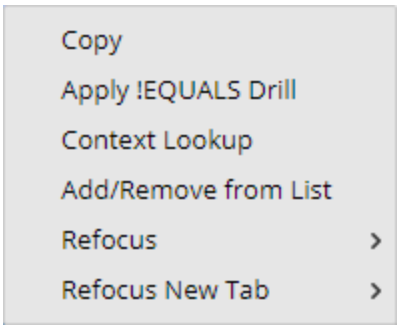
If you have the NetWitness Endpoint thick client application installed, you can launch it through the context tooltip. From there, you can further investigate a suspicious IP address, Host, or MAC address.

1. In the Events panel, the Event Header, or the Event Meta panel, hover over any underlined entity.
2. In the **ACTIONS** section of the tooltip, select **Pivot to Endpoint Thick Client**.
   The NetWitness Endpoint thick client application opens outside of your web browser.

For more information on the thick client, see the *NetWitness Endpoint User Guide*.

View the Context Lookup Panel in the Navigate View or Legacy Events View

1. Hover over different meta values to see the data sources for which data is available.
   A hover box displays a list of the data sources that have context data available for meta value. These are the possible data sources: NetWitness Endpoint, Incidents, Alerts, Hosts, Files, Feeds, and Live Connect.
2. Right-click a meta value, and click **Context Lookup** in the drop-down menu to open the Context Lookup panel.



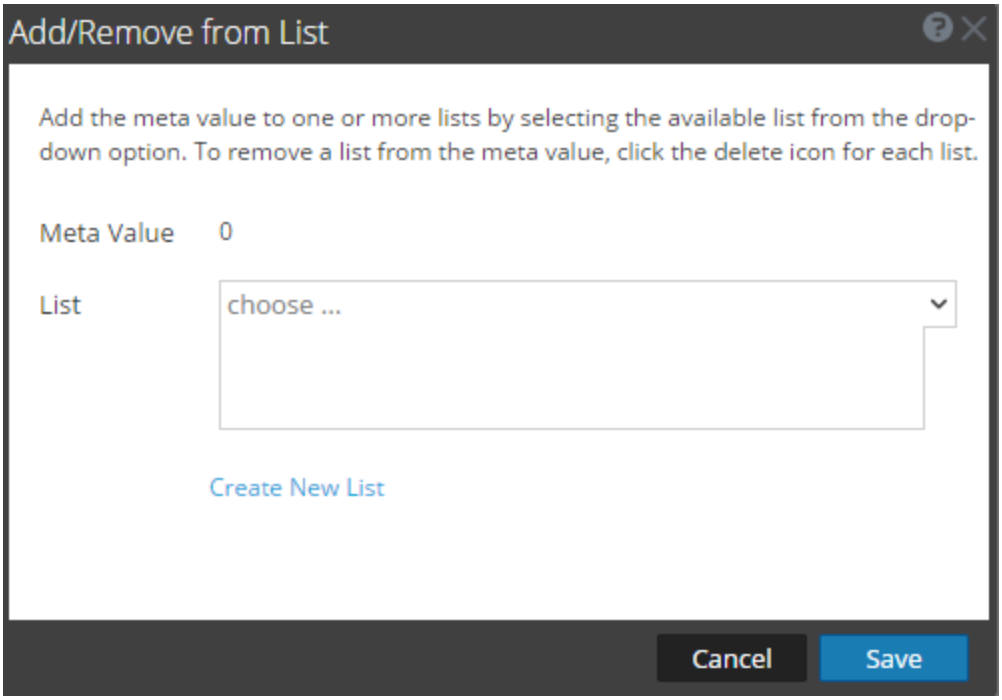
The Context Lookup panel opens from the right side of the browser window. The Context Lookup panel is populated with the information from the Context Hub as it becomes available.

3. To perform actions from the Context Lookup panel, right-click an entity such as IP address.
   The following options are available: Open Link in New Tab, Query in Investigate, Copy Link, Paste, Google Lookup, Virus Total Lookup, and Query in Endpoint.
4. To close the Context Lookup panel, click ☒ in the panel.

Add Meta Values to an Existing List (Navigate and Legacy Events Views)

To add a meta value to an existing list in Context Hub:

1. While investigating a service in the **Navigate** view or the **Legacy Events** view, right-click a meta value (for example, values under Source IP, Destination IP, or Username) and select **Add/Remove from List** in the context menu.
   The Add/Remove from List dialog is displayed.



2. In the **List** field, select one or more lists from the drop-down option to which the meta value must be added.
3. Click **Save**.
   The meta value is added to the selected lists.

Remove a Meta Value from a Context Hub List (Navigate and Legacy Events Views)

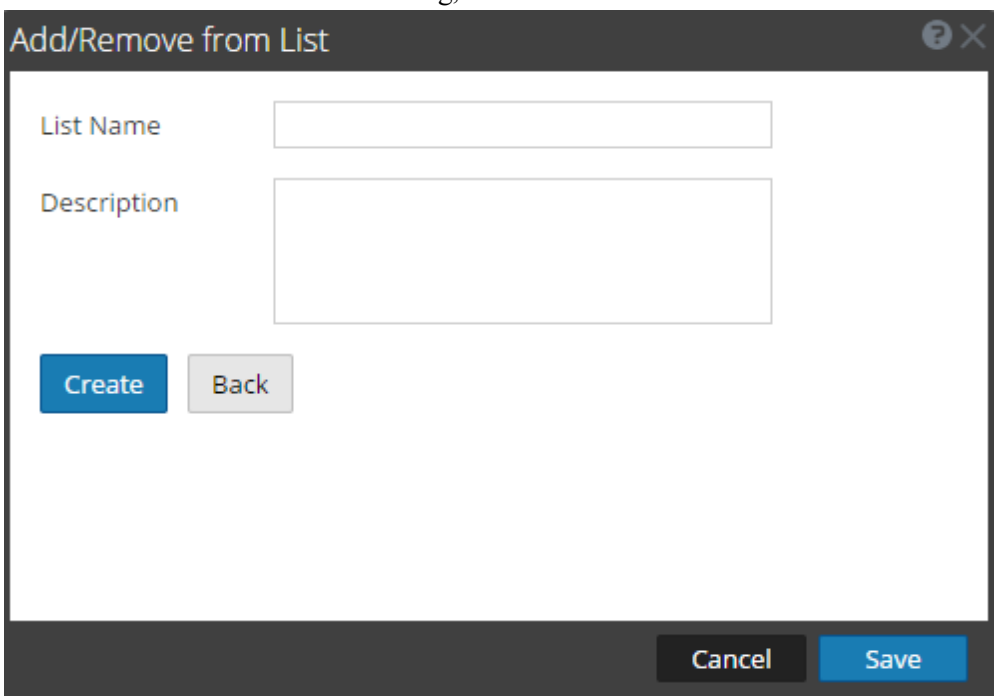
To remove a meta value from list:

1. In the **Add/Remove from List** dialog, in the **List** field, view the lists which include the meta value.
2. Click the delete icon (x) for each list that should not include the meta value.
3. Click **Save**.
   The meta value is removed from the deleted list.

Create a New List (Navigate and Legacy Events Views)

To create a Context Hub list in Investigate:

1. In the **Add/Remove from List** dialog, click **Create New List**.



2. In the **List Name** field, enter an unique name for the list.
3. In the **Description** field, enter the description of the list.
4. Click **Create** to create the list.

5. Click **Save** to add the meta value to the created list.
   These lists are considered as data sources for retrieving context information.

Launch a Lookup of a Meta Key

When you have found data of interest in the Navigate view, the Events view, or the Legacy Events view, you can do internal lookups to NetWitness Endpoint and RSA Live, as well as external lookups of meta values in community resources such as SANS IP History and ThreatExpert Search.

Analysts can use the external lookups to save time during investigations. The external lookups are available by right-clicking one of the these meta keys: IP address (`ip-src, ip-dst, ipv6-src, ipv6-dst, orig_ip`), `host (alias-host, domain.dst)`, `client`, and `file-hash`.

For all `ip` and `host` meta keys, the following lookups are built in to NetWitness:

- Google Malware: Opens a Google Malware search in a new tab.
- SANS IP History: Opens a SANS IP History search in a new tab.
- McAfee SiteAdvisor: Opens a McAfee SiteAdvisor search in a new tab.
- Endpoint Thick Client Lookup: Opens a search in the NetWitness Endpoint Thick Client in a new tab.
- BFK Passive DNS Collection: Opens a BFK Passive DNS collection search in a new tab.
- CentralOps Whois for IPs and Hostnames: Opens a CentralOps Whois search for IPs and hostnames in a new tab.
- Malwaredomainlist.com Search: Opens a Malwaredomainlist.com search in a new tab
- Robtex IP Search: Opens a RobtexIP search in a new tab.
- ThreatExpert Search: Opens a ThreatExpert search in a new tab
- IPVoidSearch: Opens a UrlVoid Search in a new tab n a new tab

For the `file-hash` and `alias-host` meta keys, the Google lookup opens a Google search in a new tab.

For the `client` meta key, the NetWitness Endpoint Lookup option opens an Endpoint Thick Client in a new tab if the client is installed on the same system on which the browser is being used.

Administrators can add additional external lookups and other custom actions as described in "Add Custom Context Menu Actions" in the *System Configuration Guide*.

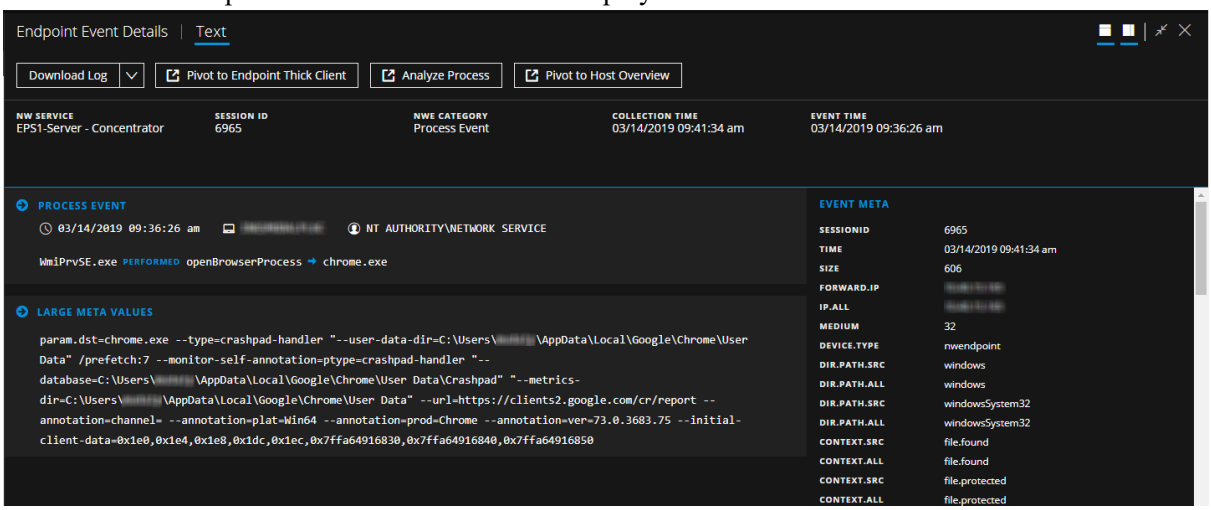
Launch an Endpoint Thick Client Lookup in the Events View

When viewing an endpoint event in the Text panel, you can pivot to analyze the same event in NetWitness Endpoint.

Note: Version 4.4.0.x of the NetWitness Endpoint (NWE) thick client must be installed on the same server, the NWE meta keys must exist in the `table-map.xml` file on the Log Decoder, and the NWE meta keys must exist in the `index-concentrator-custom.xml` file. The NWE thick client is a Windows only application. Complete setup instructions are provided in the *NetWitness Endpoint User Guide* for Version 4.4.

To open an event in NetWitness Endpoint:

1. Starting from the Navigate view:
   a. In the **Query** drop-down, select **Advanced**, and enter one of the following queries: `nwe.callback_id exists` or `device.type='nwendpoint'`
   Endpoint data is displayed in the Values panel.
   b. Right-click an event, and select **Events** in the menu.
2. (Version 11.1 and later) Go to **INVESTIGATE > Events**. In the **Query** drop-down, select **Advanced**, and enter one of the following queries: `nwe.callback_id exists` or `device.type='nwendpoint'`
   Endpoint data is displayed in the Events panel.
3. Select an event.
   The Events view opens with the selected event displayed in the Text view.



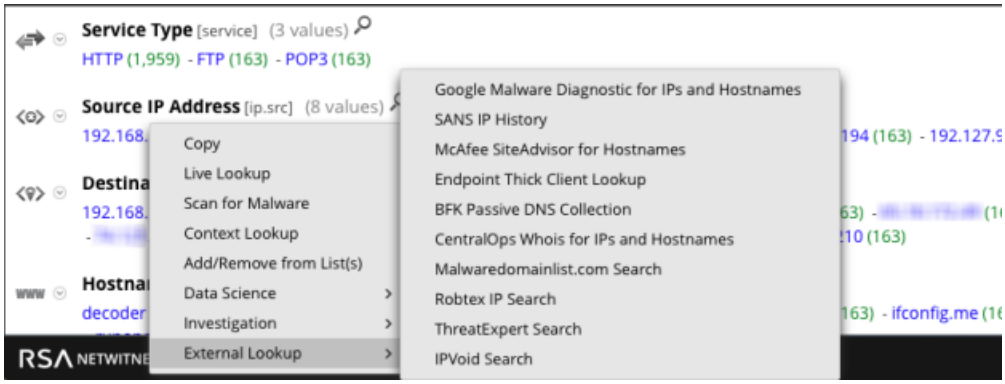
4. In the Event Header click **Pivot to Endpoint.**
   A new browser tab with the url `ecatui://<id>` opens and the NWE Thick Client is launched. If the NetWitness Endpoint Thick Client is not installed, no data is displayed and the following message is displayed: `Applicable for hosts with 4.x Endpoint agents installed, please install the NetWitness Endpoint Thick Client.`

Launch an Endpoint Thick Client Lookup in the Navigate View

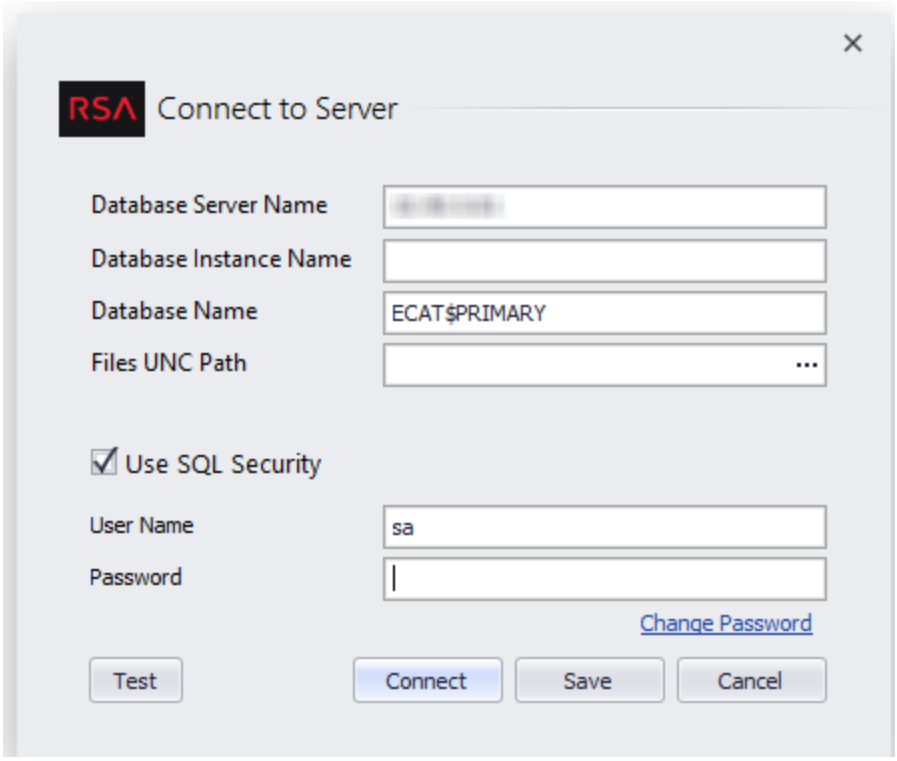
To launch an Endpoint Thick Client lookup of data from the Navigate view:

1. Right-click a meta value for one of the following meta keys: `ip-src`, `ip-dst`, `ipv6-src`, `ipv6-dst`, `orig_ip`, `alias-host`, `domain.dst`, or `client`.

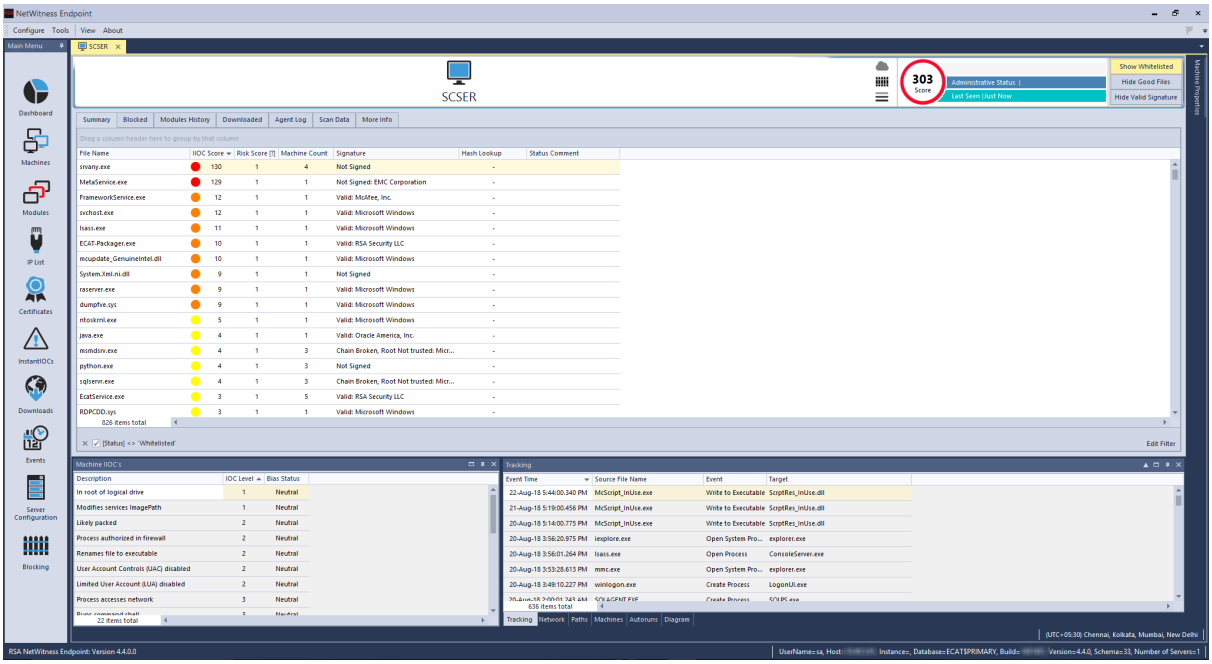
2. Select **External Lookup** in the context menu.
   A submenu of external lookup options is displayed.



3. Select **Endpoint Thick Client Lookup**.
   The Connect to Server dialog is displayed.



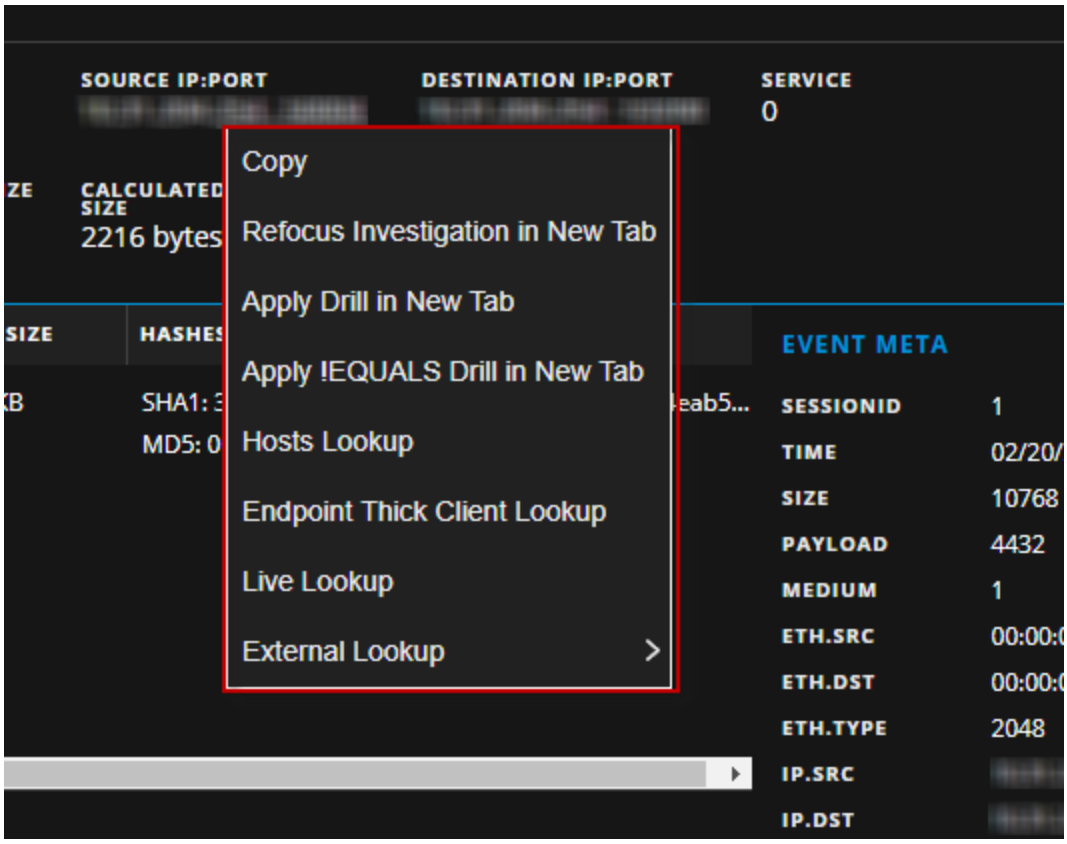
4. Enter the user name and password required to log in to the Endpoint Thick Client, and click **Connect.** The drill point opens in NetWitness Endpoint.



Perform Lookups of Meta Values in Events

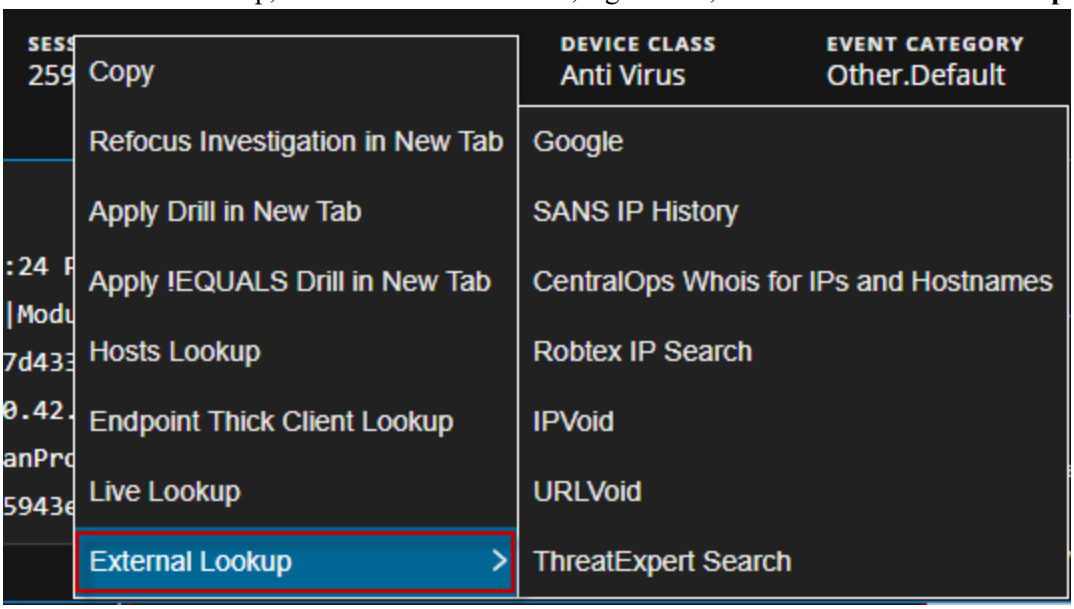
In the Events view, you can further investigate meta values in an event by right-clicking certain meta values and using the options in a drop-down menu. Not all fields have right-click actions. To perform internal and external lookups:

1. In the Events view, right-click a meta value in the Events List, the Event Meta panel, or the Event Header. Some meta values have a drop-down menu.



2. Select one of the following internal lookups:
   - **Copy**: Copies the meta value to the clipboard.
   - **Refocus Investigation in New tab**: Launches the another investigation in a new tab with the focus on the selected meta value.
   - **Apply Drill in New Tab**: Applies the drill and launches it in a new tab to drill the data in Navigate view.
   - **Apply !EQUALS Drill in New Tab**: Applies (!EQUALS) to the meta and launches a new tab, effectively excluding the meta value from the results.
   - **Hosts Lookup**: Looks up the value in the Investigate > Hosts view.
   - **Endpoint Thick Client Lookup**: Analyzes the meta value in the Endpoint Thick Client (for clients which have Endpoint Agent).
   - **Live Lookup**: Looks up a meta value on Live for further analysis.

3. For an external lookup, hover over a meta value, right-click, and select **External Lookup**.

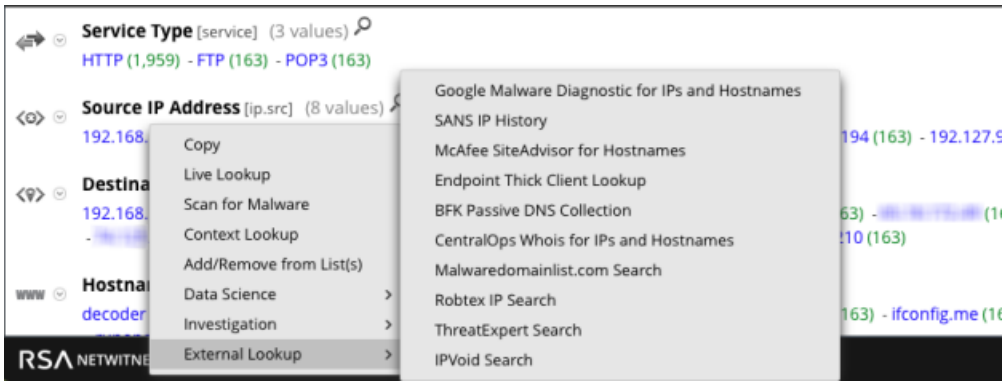


4. In the submenu select one of the available external lookups:
   - **Google**: Looks up a meta value on Google.com
   - **SANS IP History**: Looks up a meta value on SANS IP History, domain = http://isc.sans.org/ipinfo.html?ip=ipaddress
   - **CentralOps Whois for IPs and Hostnames**: Looks up a meta value on CentralOps Whois for IPs and Hostnames, domain = http://centralops.net/co/DomainDossier.aspx?addr=domain&dom_whois=true&dom_dns=true&net_whois=true
   - **Robtex IP Search**: Looks up a meta value on Robtext IP Search, domain = https://www.robtex.com/cidr/domain.ipaddress
   - **IPVoid**: Looks up a meta value on IPVoid, domain = http://www.ipvoid.com/scan/domain/
   - **URLVoid**: Looks up a meta value on URLVoid, domain = http://www.urlvoid.com/scan/ipaddress/
   - **ThreatExpert Search**: Looks up an IP meta value on ThreatExpert Search, domain = http://www.threatexpert.com/reports.aspx?find=IP address

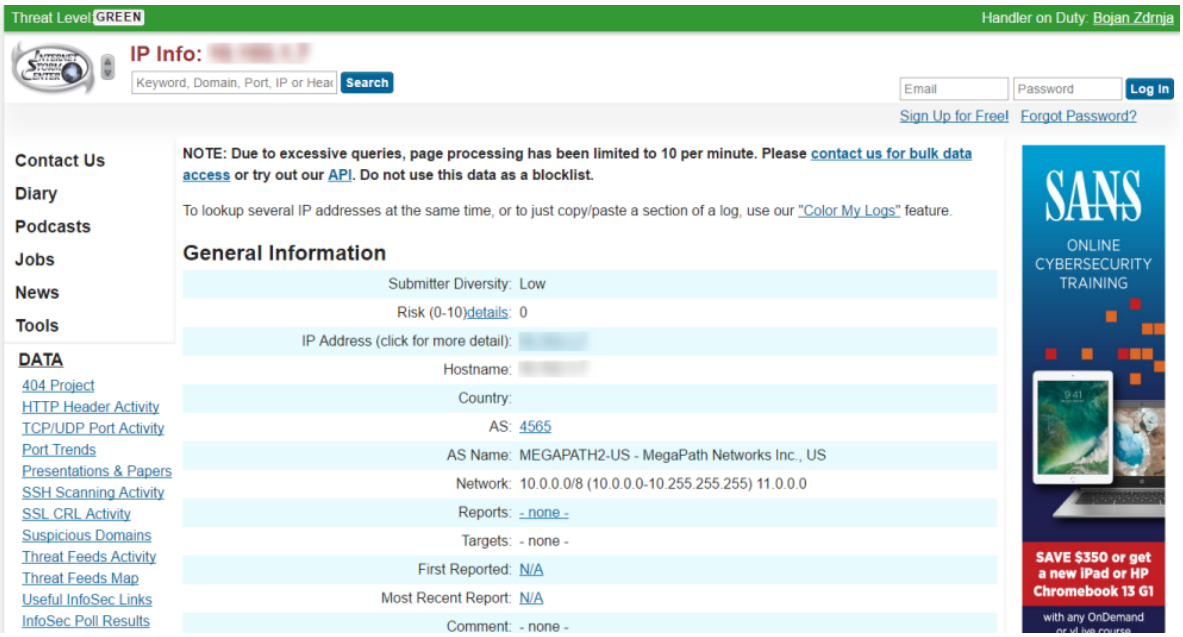
Launch Other External Lookups from the Navigate View

To launch an external lookup of data from the Navigate view (other than NetWitness Endpoint Thick Client Lookup):

1. Right-click a meta value for one of the following meta keys: `ip-src`, `ip-dst`, `ipv6-src`, `ipv6-dst`, `orig_ip`, `alias-host`, `domain.dst`, or `client`.
2. Select **External Lookup** in the context menu.
   A submenu of external lookup options is displayed.



3. Select one of the lookup options.
   The selected meta value opens in the selected lookup, for example, if you selected SANS IP History, the drill point information is displayed in SANS Internet Storm Center.



Launch a Malware Analysis Scan from the Navigate View

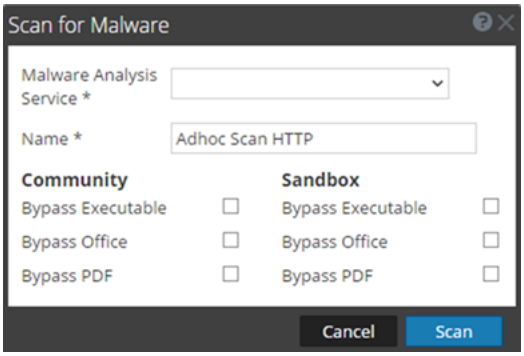
From within Investigate, analysts can launch an on-demand Malware Analysis scan by selecting a service and meta value, and choosing an option from the context menu. When polling is complete, the scanned data is available for malware analysis.

To launch a Malware Analysis scan of data from the INVESTIGATE > Navigate view:

1. Right-click a meta value (for example, OTHER, DNS, or FTP) and select **Scan for Malware** in the context menu.

   The Scan for Malware dialog is displayed with a suggested name for the on-demand scan and no service selected.

2. In the Scan for Malware dialog, select a service to perform the scan, edit the name, and select the types of files to bypass under community and sandbox.

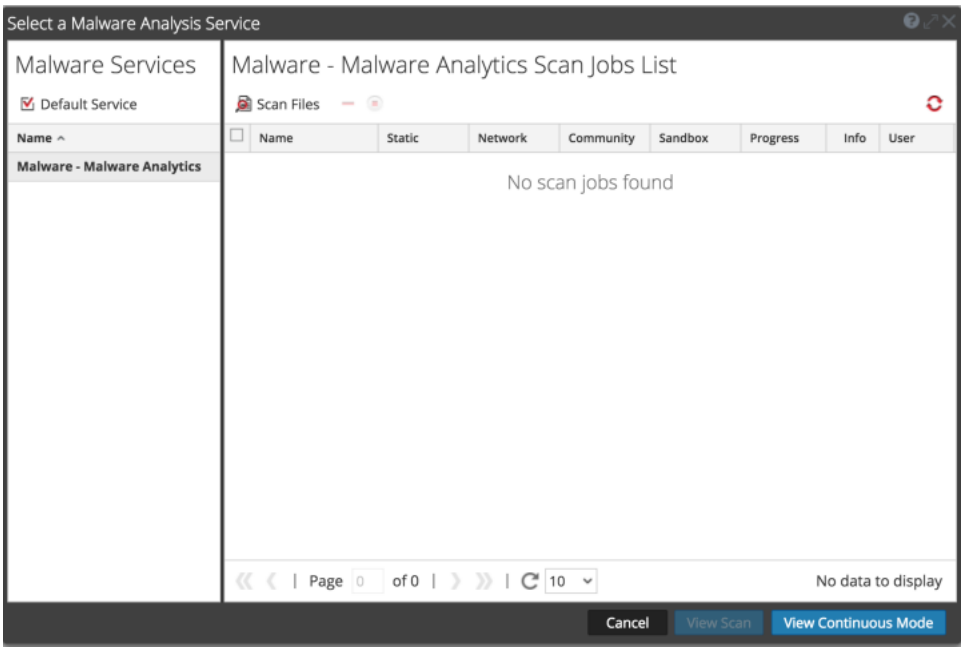


3. Click **Scan**.

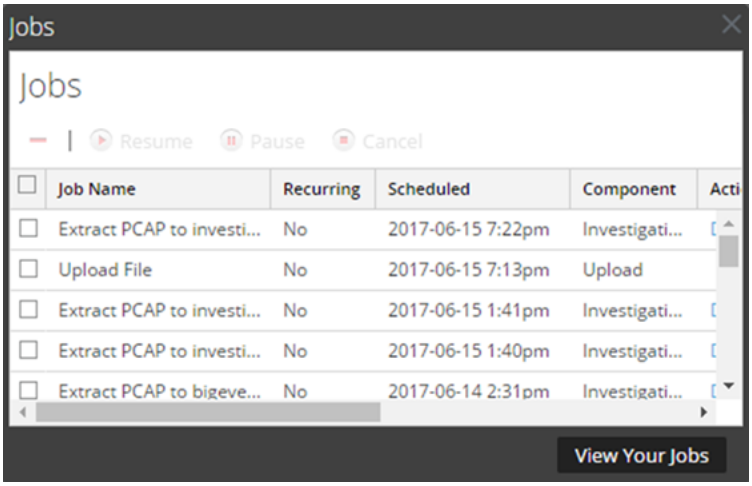
   The scan request is added to the Scan Jobs List dashlet and the Jobs Tray. The bypass settings in this dialog override the default settings in the basic Malware Analysis configuration settings.

4. To view the jobs, do one of the following:

   a. Go to the Scan Jobs List in the Malware Analysis view or in the Unified dashboard. Double-click a scan to view the scan.



b. To view the job in the Jobs tray, click  in the NetWitness toolbar. When the job is complete, scroll to the left and click **View**.



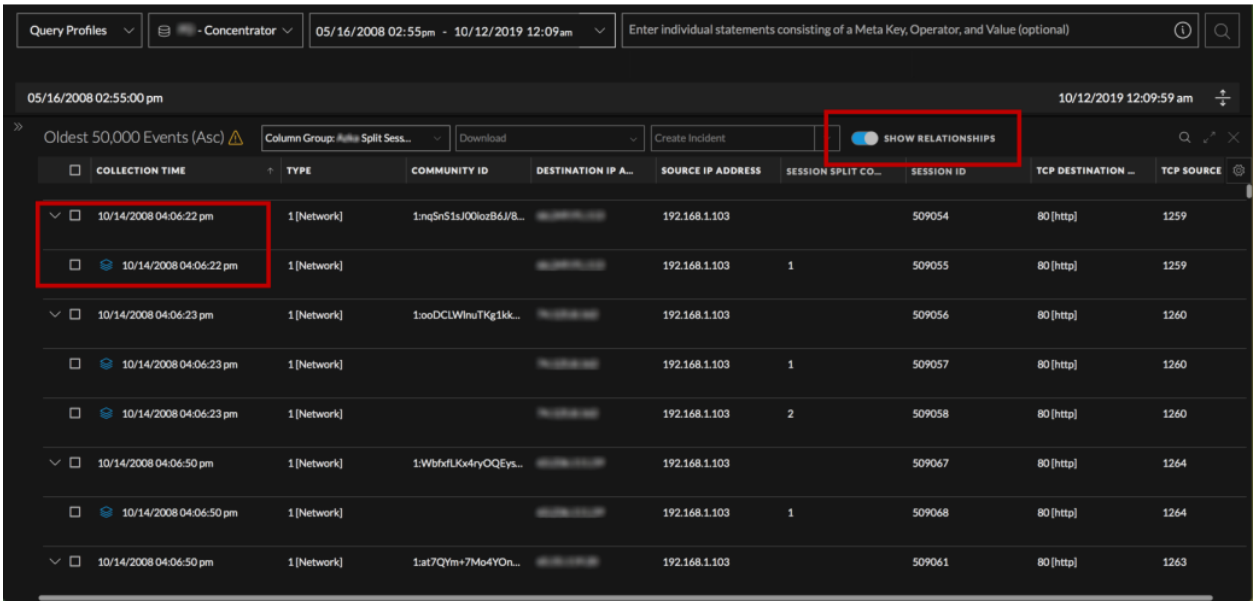
The Malware Summary of Events for the selected scan is displayed. The scan is also added to the list of available scans in the dialog for selecting scans in the Investigation > Malware tab.

Group Events from Split and Related Sessions in the Events and Legacy Events Views

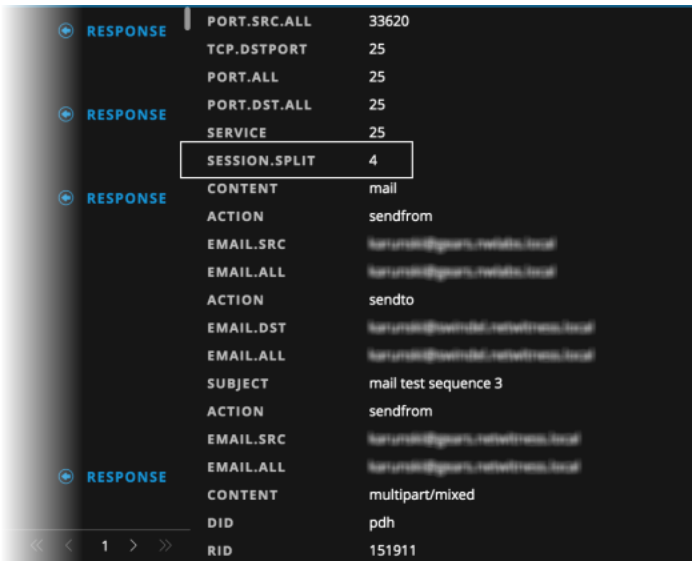
In the Events view listing of events, events from split and related sessions are listed in the order they were parsed; for this reason, they are not always listed together in the Events list. In Version 11.4.1 and later, the Group Events option allows you to change the order in which the events are listed so that you can more easily detect relationships in the captured data. When events are grouped, we refer to the first event as the leading event.

The user interface is designed to help you identify grouped events. Solid row lines delineate different groups of related events while dotted lines depict the events that are part of the same related group. In a group of events, the leading event is first and subsequent events are nested under the leading event with indentation and relationship icons for subsequent events. The numbers next to the relationship icon distinguish the session split count.

If the leading event is not included in the current data set, the subsequent events are still grouped together under the first subsequent event. Only the leading event or the first subsequent event, in the case of no leading event, is sorted; indented events are not sorted. Hovering over the subsequent event marker () displays a tool tip explaining the relationship. The following figure is an example of related events as they are displayed in the Events list.



If events are related based on session fragments, when you select a subsequent event and open the reconstruction, you can see the `session.split` meta key in the Event Meta panel.



Split Network Session

If you see a tool tip like this, the listed events are part of a split network session:
`The event is part of a split session (session.split: #) matching these parameters: ip.src=127.0.0.1 AND ip.dst=127.0.0.1 AND tcp.srcport=25 AND tcp.dstport=1234.`

The reason for the split is one of the following:

- The original event was split into sub-parts by creating additional events for each transaction inside the original event.
- The original session was split as it was ingested into the Network Decoder because the size was larger than the Assembler Maximum Size (default=32 MB).
- The original session was split as it was ingested into the Network Decoder because the time was longer than the Assembler Timeout Session (default=60 sec).

Session Size and Time Split

Network Decoders are configured with a default session size (`assembler.size.max`) and a session timeout (`assembler.timeout.session`). The configuration is described in "Configure Session Split Timeouts" in the *Decoder Configuration Guide*. When a session exceeds either limit, the Network Decoder splits the session, and all subsequent packets become part of a new session, fragmenting the actual network session across multiple Network Decoder sessions. To process session fragments with the context that they are fragments of the larger network session and to improve association of source and destination addresses, ports, and application protocols, Network Decoders parse contextual fragments and highlight session fragments.

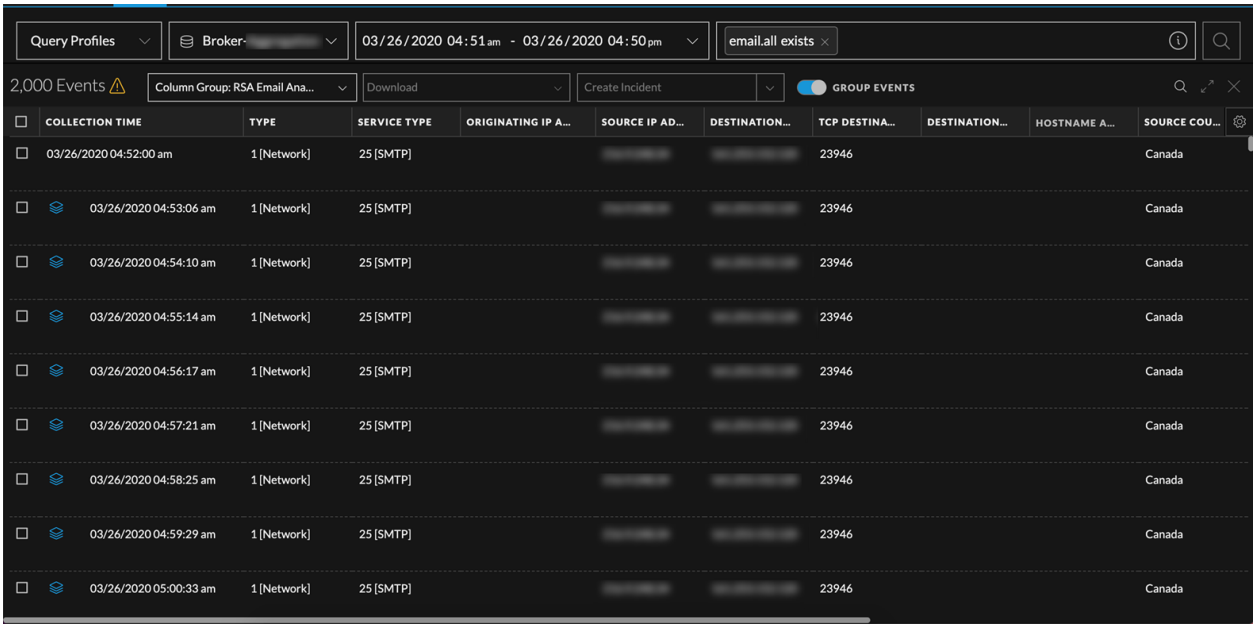
Note: In the Legacy Events view, you can find session fragments and export all packets viewed in the Events list to a single PCAP as described in Find and Combine Fragments in the Legacy Events View.

The Network Decoder completes session parsing before splitting the session based on the configured maximum session size (default = 32 MB) or the configured timeout (default = 60 seconds). When parsing is complete, the parsed results include the proper address directionality and application protocol, which are propagated to each subsequent session fragment to ensure consistency with the logical network session they represent.

Transaction Handling Split

Administrators can configure a Network Decoder to subdivide incoming sessions into smaller transaction sessions when using Lua parsers that are designed to create transactions. The configuration is described in "Configure Transaction Handling on a Decoder" in the *Decoder Configuration Guide*. The Decoder service configuration node has a parameter that controls the behavior of the Network Decoder when a parser defines a transaction within a network session: `/decoder/parsers/config/parser.transaction.mode`. If the mode is set to `split`, a large session with multiple application-level transactions is split when a parser generates an application-level transaction such as an email. An example of this is a large session containing multiple emails. For each email (transaction), a new session item (split session) is created, network meta items are copied into the new session, and meta items marked in the transaction are copied from the original session to the new session.

Transactions require a parser update to function, and initially they only support SMTP and HTTP pipelining use cases. This is an example of the reconstruction of an email that is separated based on the individual emails inside the original event. Each transaction highlights a single email and the metadata associated with the transaction is only related to that email. To provide this functionality, the original packets are stored on the Network Decoder as usual for a network event, but the new related transaction events are created on the Concentrator. As a result, analysts see visual queues in the user interface, and they also have the ability to query to find only specific emails or attributes of emails that were previously all bundled together. To eliminate the original event from the query results, the `session.split` meta key has been indexed. When there are transaction splits, the original event does not have that meta key associated with it while all the related transactional events do.



Session Fragments Highlighting

Both types of session fragments have an additional meta key: `session.split`. The earliest fragment is session 0 and sessions with a later time stamp are incrementally numbered 1, 2, 3, and so on. The `session.split` meta key indicates the number of preceding sessions fragments; however, it does not always indicate that there are subsequent session fragments, even with a value of 0. It is also possible for the first fragment of the session to not have `session.split` metadata if the session is parsed before exceeding the maximum session size.

Transaction splits start with a `session.split` value of 1. When viewing sessions, the `session.split` meta key clearly identifies sessions that are fragments in the Events view and the Legacy Events view (Events List view and the Events Detail view).

If this was a session size and timeout split, you can view the session fragments and determine the maximum session size or session timeout necessary for parsing to combine the split sessions into one again. For example, if you have four fragments at 32 MB, you need to configure your test Decoder (usually a virtual machine set up separately from the production service) with a maximum session size greater than 128 MB. The steps are the same to find all fragments based on a session timeout.

Related Network Session

If you see a tool tip like this, separate events processed by the Network Decoder share a set of four values that identify the IP source, IP destination, source port, and destination port:

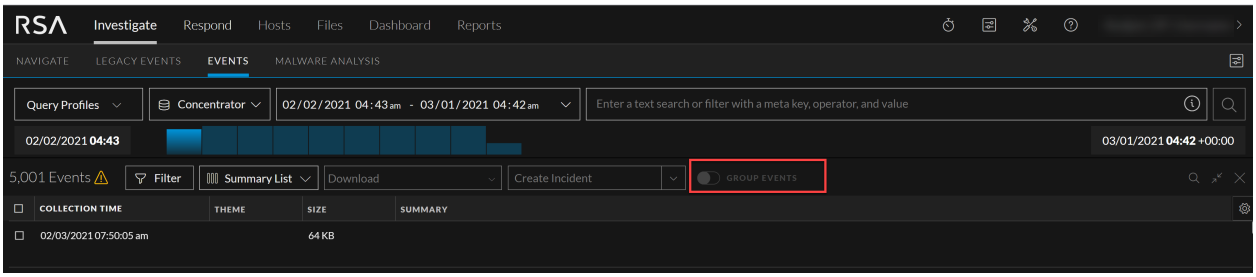
```
The event is related to a previous session matching these parameters:
ip.src=127.0.0.1 AND ip.dst=127.0.0.1 AND tcp.srcport=25 AND
tcp.dstport=1234"Second category: Related Network Session
```

In this case, the Network Decoder has not inserted a split, and there is no `session.split` metadata associated to any of the events. The reason these events grouped together is to highlight the events that are worthy of scrutiny based on a pattern. Each event has the same source IP address, the same destination IP address, the same source port, and the same destination port. The grouping of related events does not occur if any of the four meta keys is obfuscated for data privacy reasons.

These are the combinations of meta keys that must match to categorize an event as a Related Network Session:

- `ip.dst`, `ip.src`, `tcp.dstport`, `tcp.srcport`
- `ip.dst`, `ip.src`, `upd.dstport`, `udp.srcport`
- `ipv6.dst`, `ipv6.src`, `tcp.dstport`, `tcp.srcport`
- `ipv6.dst`, `ipv6.src`, `upd.dstport`, `udp.srcport`

In Version 11.6, the group events option that retrieves events from logs, network and endpoints is disabled when data privacy is enabled. For example, if the `'ip.src'` is blacklisted or restricted by the admin, then the option shown in the following image is disabled.



This option is disabled if one or more of the following 12 metas are restricted:

`'sessionid'`

`'nwe.callback_id'`

`'medium'`

`'session.split'`

`'ip.dst'`

`'ip.src'`

`'ipv6.src'`

`'ipv6.dst'`

`'tcp.dstport'`

`'tcp.srcport'`

`'udp.dstport',`

`'udp.srcport'`

Use Cases for Viewing Events from Split and Related Network Sessions

These are examples of practical use cases for viewing events from split sessions:

- A Network Decoder that is inline with a proxy server receives a lot of email connections that get bundled into a single session based on event time as NetWitness sees them. The single session has multiple meta values for `subject`, `email.src`, `email.dst`, and other meta keys relevant to email, which are difficult to map together correctly. Seeing the session organized as leading and subsequent events, gives the analyst a clear idea of the details of each email.
- An analyst is attempting to understand what IP address out of all the metadata associated with a session resulted in the generated metadata or an alert, but the IP address is not in the output. For example, a feed that is parsing for an indicator of compromise may have many triggers in a session that has many IP addresses. By viewing the complete event organized as leading and subsequent events, the analyst can understand which IP triggered the alert.
- An analyst wants to know which file was deleted from which directory versus which file was read out in which directory, but the session encompasses multiple files and directories. For example, an HTTP connection with commands: `directory /keep/`, `directory /temp/`, `filename foo.txt`, `filename me.doc`, `action delete`, and `action read`. Viewing the leading and subsequent

events tells the analyst that `/temp/me.doc` was deleted and `/keep/foo.txt` was read. Now the analyst or analytics can form an opinion about how bad these actions really were.

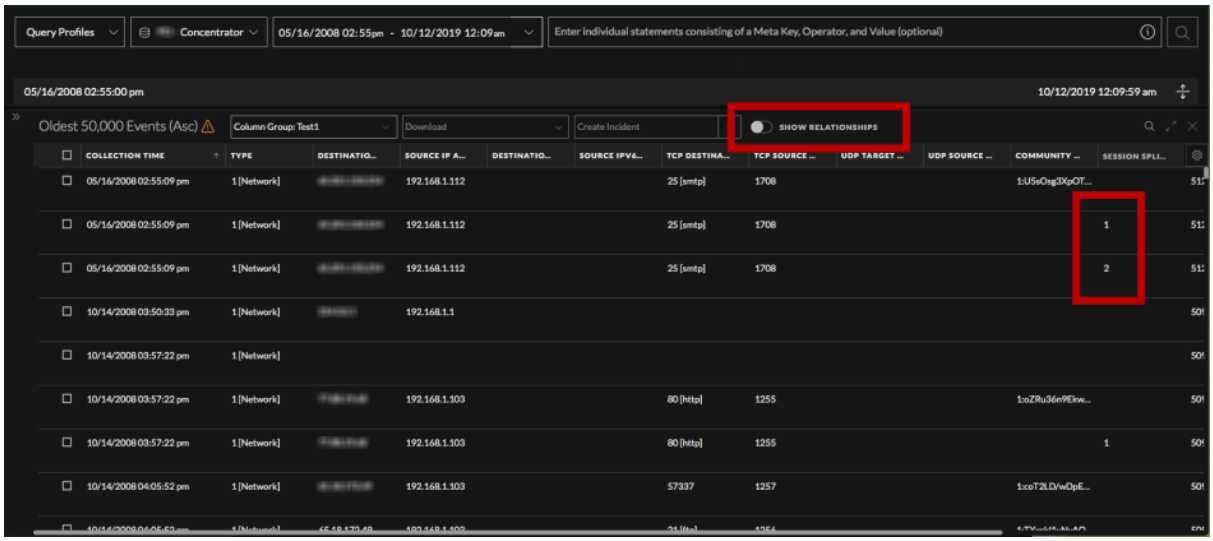
- An analyst is attempting to retrieve a large file related to an event that triggered a suspicous alert. However, the file that was transferred was so large that the network decoder split it into 100 separate sessions. When viewing this group related split sessions, the analyst can download a PCAP of the sessions, then extract the original file by either running it through a Decoder with larger assembler settings or a third-party tool.

Show and Hide Relationships in the Events List

For both types of related events, you can see the relationship of events in the Events view Events list. When the Events list is first displayed, you can tell if set of results includes related events by looking at the Group Events switch at the top of the Events list. If the results do not include related events, the switch is dimmed as shown in the following figure.

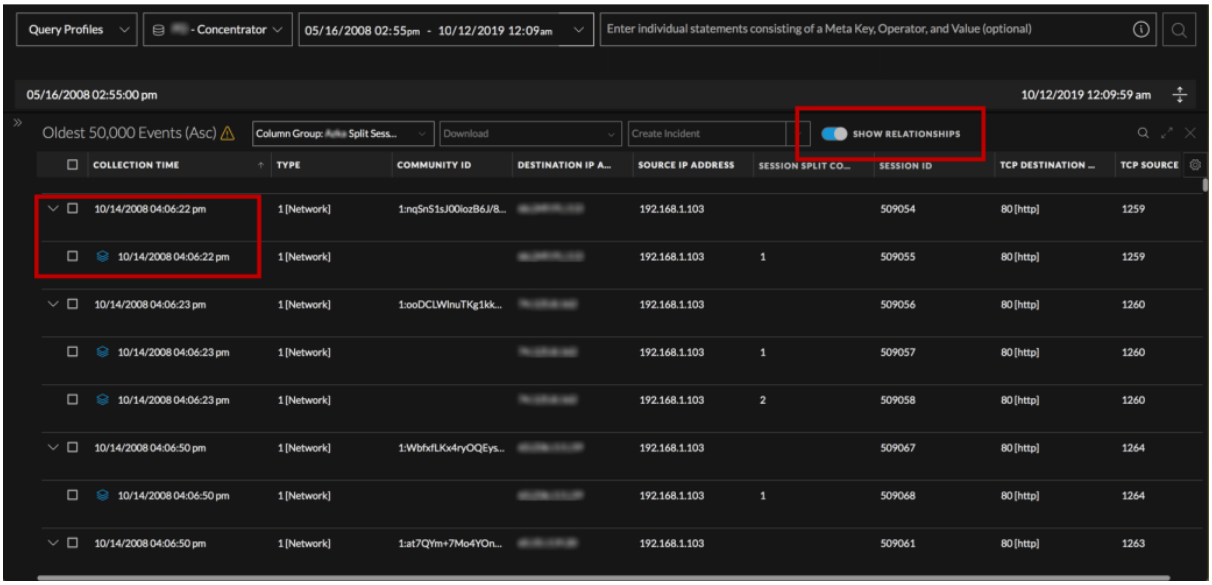
To look for related events in the Events list:

1. Go to **Investigate > Events** and submit a query.
   If the results include related events, the Group Events switch is active, but not enabled. The figure below illustrates a set of results that include split sessions, and the Group Events switch is disabled. The related events are not nested.



2. Click the **Group Events** switch.
   Related subsequent events are nested below the leading event. The subsequent events are indented

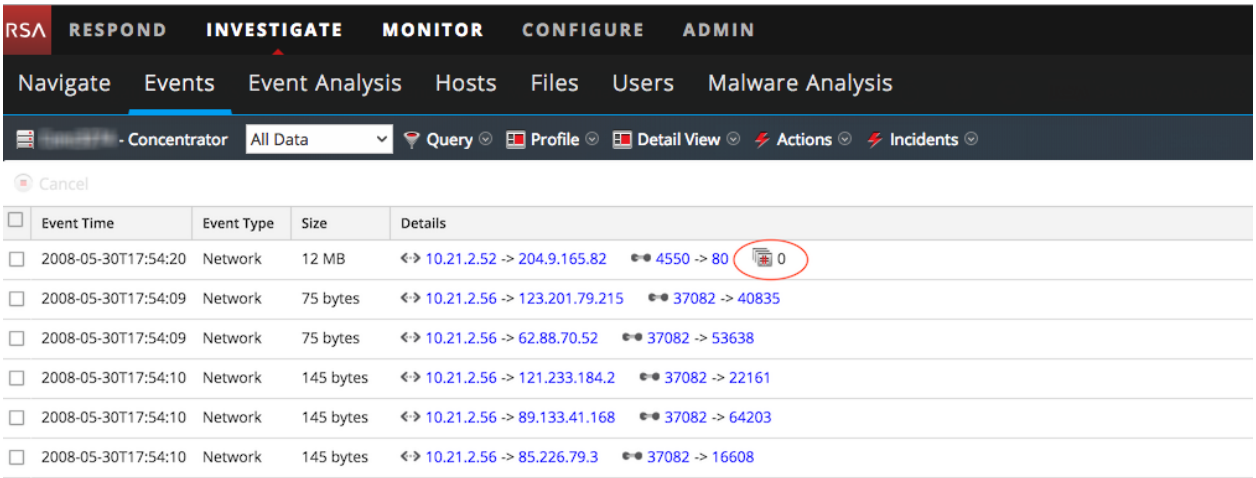
and marked by an icon. Clicking on the icon explains why the event is grouped.

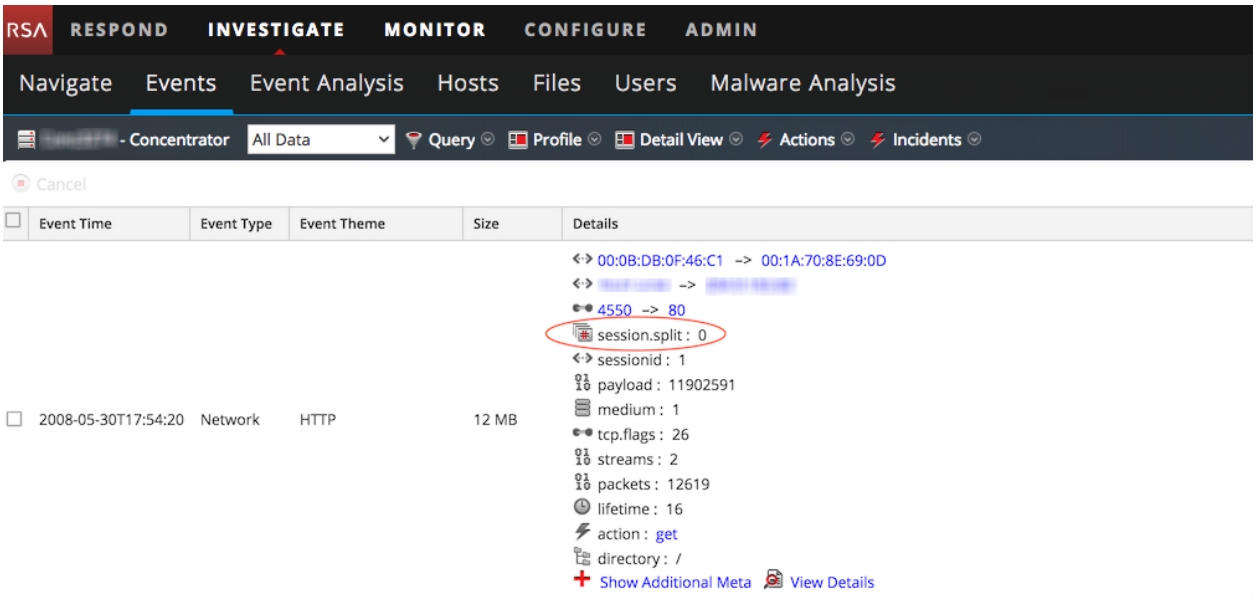


Session Fragments Highlighting (11.3 Events View)

The figures below show the Events List view and the Events Detail view with fragmented session information highlighted.

Note: A maximum session size of 12 MB was configured at the time the screen captures below were created.





Note: The `session.split` meta key must be indexed to help you find the split sessions in the results.

The `session.split` metadata is displayed immediately following the address and port metadata in the details view (Legacy Events view). These enhancements make it possible to quickly:

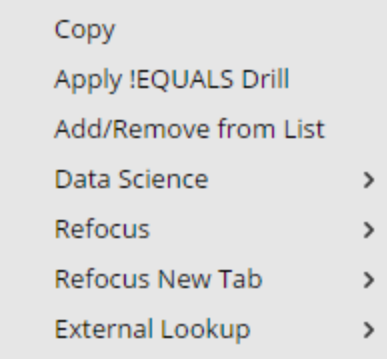
- Identify sessions that are fragments of a network sessions.
- View all of the session fragments of a network session given a single session fragment.
- Export the packets for the entire network session as a single PCAP file.

Find and Combine Fragments in the Legacy Events View

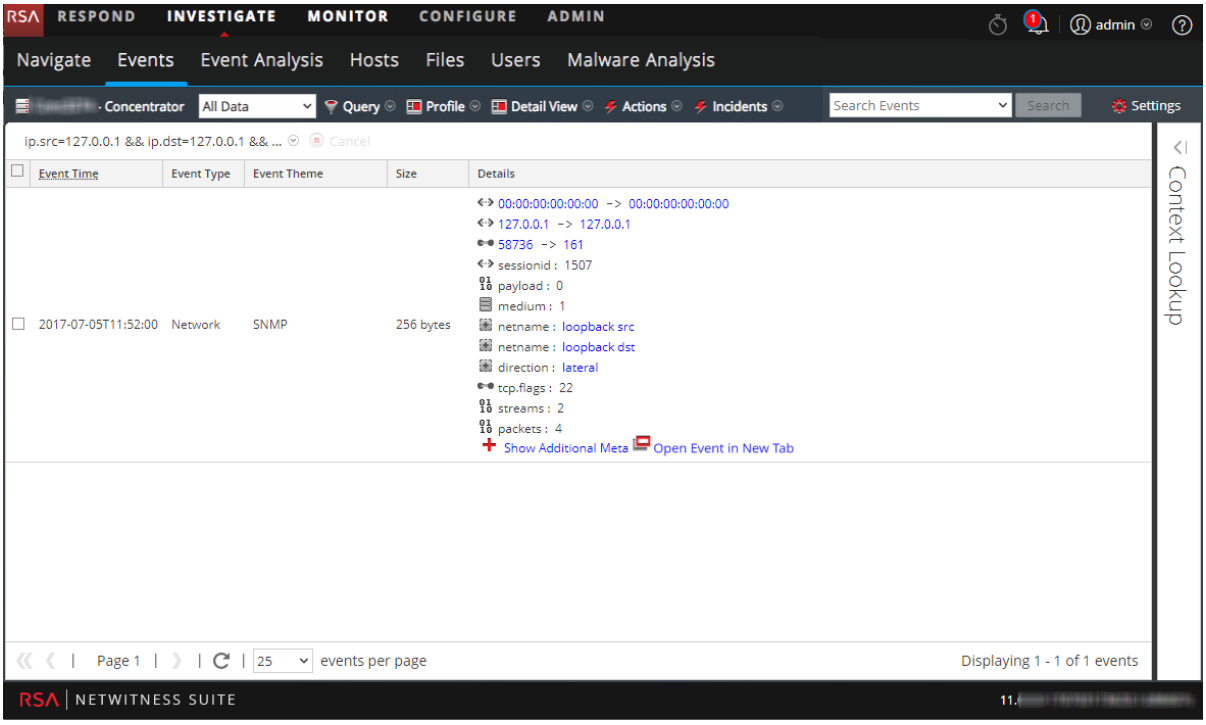
From within the Legacy Events view, you can find fragments of a session using the Refocus > Find Session Fragments context menu option. NetWitness composes a query using the source and destination addresses and ports of the selected session and displays all sessions that match that query within the current time window.

To find session fragments:

1. In the **Legacy Events** view, right-click any of the source and destination address and port values: `ip.src`, `ip.dst`, `ipv6.src`, `ipv6.dst`, `tcp.srcport`, `tcp.dstport`, `udp.srcport`, and `udp.dstport`) as well as `session.split` values.
   The context menu is displayed.



2. Select **Refocus > Find Session Fragments** or **Refocus New Tab > Find Session Fragments**. NetWitness repopulates the Events list with session fragments for a single session within the current time range. Depending on the option you selected, the refocus replaces the current view or opens in a new tab. (All data is used in these examples but is not recommended on production systems).



3. If necessary, adjust the time range to include any session fragments that may precede or follow the current time window. You can tell that the time range needs to be expanded if the fragments occur near the time boundary, especially if the first visible fragment does not have a split value of 0 (or none). Alternately, inspecting the packets of the last visible session may lead you to believe that the session continues. Here is an example:

   a. If you are looking at fragments that are obviously not the first fragment, for example, 1, 2, 3, and 4 in time range 10:30 to 10:35, there should be a fragment 0. You can increase the time range to start earlier (for this example, 10:25) to find the additional fragment.

b. If the session size of last fragment is close to maximum session size (12 MB in this example), look for additional fragments by increasing the time window to include a later time (for this example, 10:40).
When all of the session fragments of a network session are included within a single Events list, the list can span multiple pages.

4. (Optional) To export the packets for every session fragment to a single PCAP file, select **Actions > Export All PCAP**.
A message informs you that the PCAP is being downloaded. When download is complete, PCAP file includes the entire network session that was fragmented.

Visualize Metadata as Parallel Coordinates

Analysts can use the parallel coordinates visualization in the Navigate view to focus the investigation on combinations of meta keys and meta entities, and values that may indicate events are abnormal and worth investigation. The parallel coordinates chart is a way of visualizing the current drill point in Investigate to examine more than two meta keys simultaneously. Visualizing multiple meta keys simultaneously can help in identifying security issues associated with multivariate patterns and comparisons, such as when individual meta keys and values may not be of concern, but combining them together may bring an abnormal pattern or relationship to light. Meta groups (see Use Meta Groups to Focus on Relevant Meta Keys) can be used effectively to define a collection of meta keys that you want to visualize as parallel coordinates.

Best Practices for Effective Parallel Coordinates Charts

To create effective parallel coordinates charts, follow these recommendations:

- Use the RSA built-in meta groups that are included in a new installation.
- Start from a drill point rather than attempting to visualize all data.
- Limit the time range if necessary.
- Choose the smallest useful set of meta keys to display as axes.
- Specify the sequence of axes to highlight anomalies between the meta values as you follow a line across the chart.
- When you can identify a useful set of meta keys and sequence, create a custom meta group to use for future investigations. For example, you can create a custom meta group for Windows executable file types.
- Reuse and share custom meta groups by importing and exporting groups as `.jsn` files.
- It may be useful to create two versions of each custom meta group. One for analysis of meta values and one for creating a parallel coordinates chart focusing on a smaller subset of the same use case.

Note: When you import meta groups, an error message is displayed if any group is already present. To import a group that is a duplicate, you must first delete the existing group. If you want to delete a meta group, it cannot be in use by a profile.

To help build better parallel coordinates charts, several optimizations are included in NetWitness.

- Analysts can specify that only sessions in which all meta keys exist are rendered in the chart.
- The administrator can increase the number of meta values rendered in the Parallel Coordinates Settings in the ADMIN > System > Investigation panel > Navigate tab.

RSA Meta Groups for Parallel Coordinates Use Cases

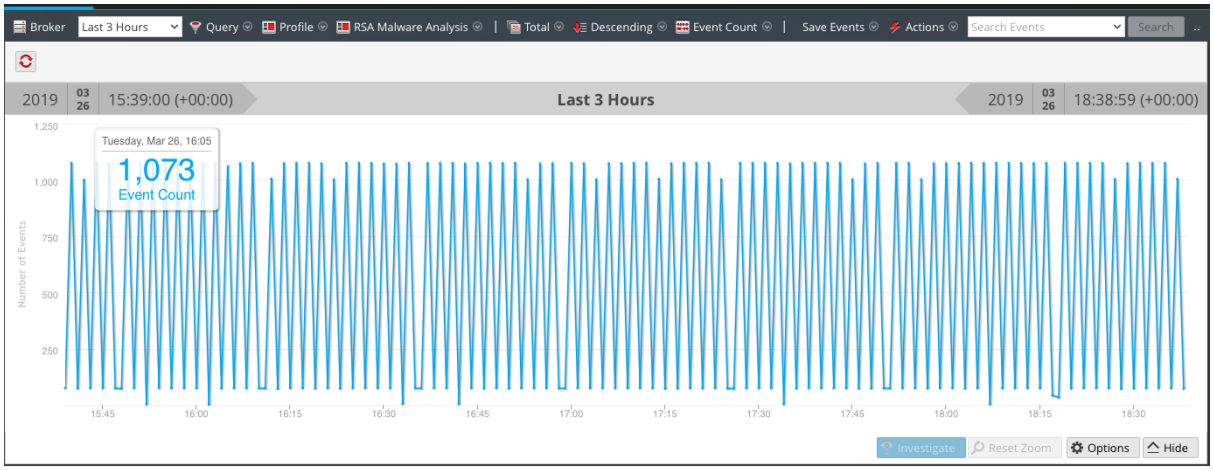
A set of predefined meta groups is included with NetWitness. If you want to get the latest version, you can import the meta groups file, `MetaGroups_ootb_w_query.jsn`, in the Manage Meta Groups dialog. Some of the targeted activities that lend themselves well to Parallel Coordinates visualizations are:

- Botnet Beaconing
- Covert Channels
- Email Analysis
- Encrypted Sessions
- Endpoint Analysis
- File Analysis
- Malware Analysis
- HTTP
- SSL/TLS
- SQL Injection Attacks
- Threat Analysis
- Web Analysis

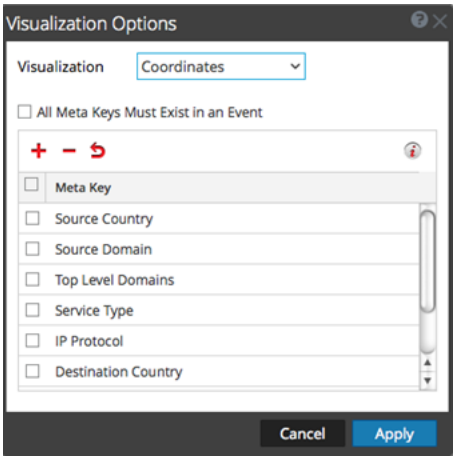
View a Parallel Coordinates Visualization

From an investigation in the Investigate > Navigate view:

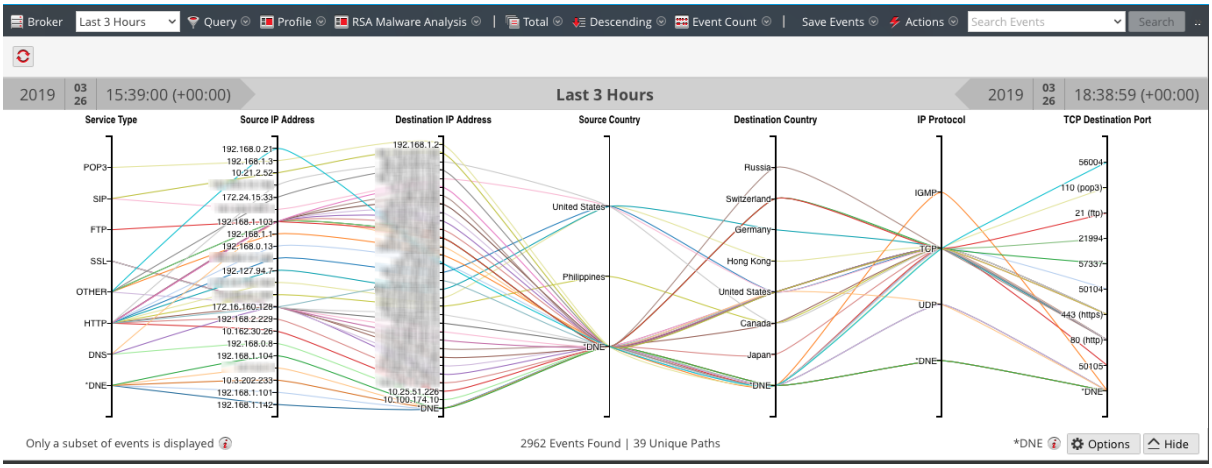
1. If the Visualization panel above the Values panel is closed, select **Visualization**.
2. In the toolbar, select **Meta** > **Use Meta Group** > **RSA Malware Analysis**.
3. A default visualization for the current drill point is displayed as a timeline.



4. In the **Visualization** panel, select **Options**.
   The Visualization Options dialog is displayed.
5. In the **Visualization** drop-down list, select **Coordinates** and click **Apply**.



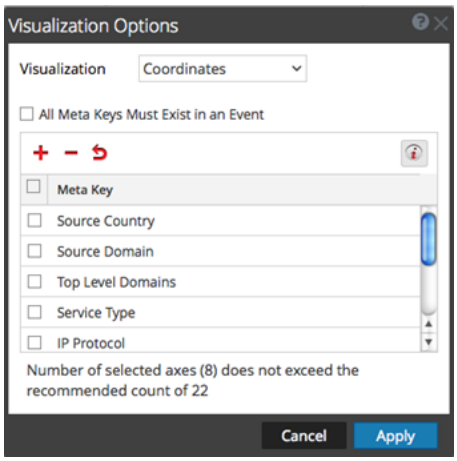
The visualization is loaded. In this example, 2962 events are found and 39 unique paths are visualized.



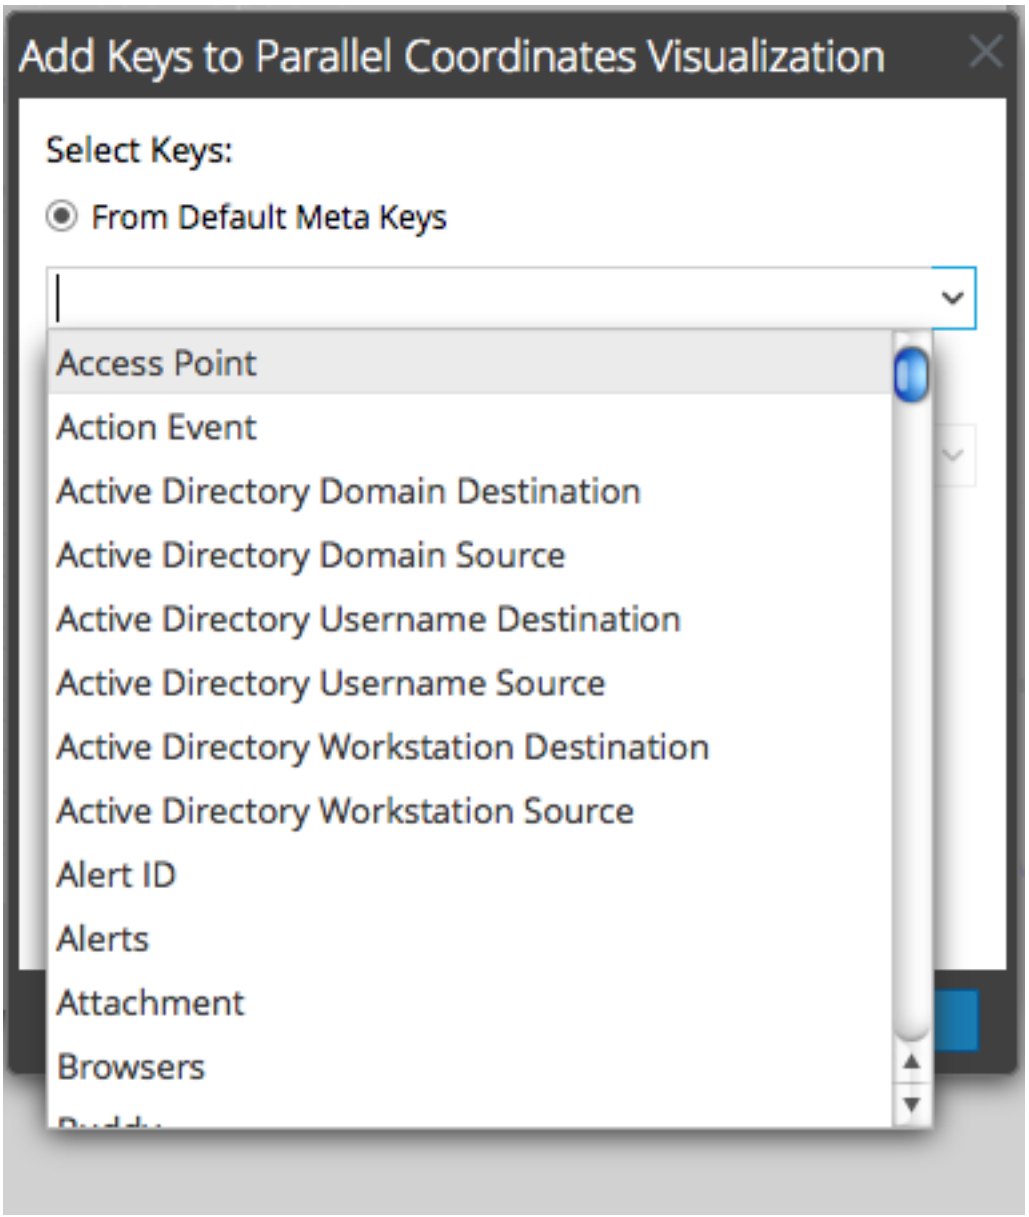
Select Meta Keys for a Parallel Coordinates Visualization

With a Parallel Coordinates visualization open, do the following:

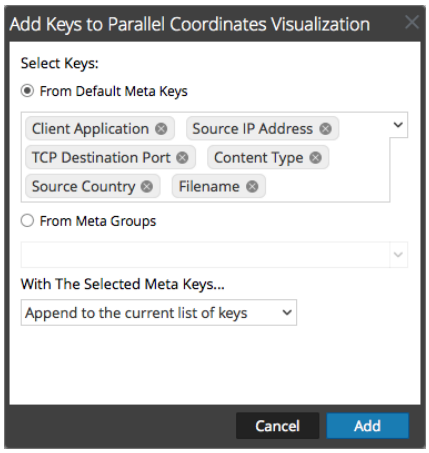
1. In the Visualization panel, select **Options**.
   The Visualization Options dialog is displayed. In the toolbar, click  to display the recommended number of axes for a readable visualization. When a recommended count of keys is displayed, the count changes based on the browser size. If you make the browser window larger, the recommended count is increased.



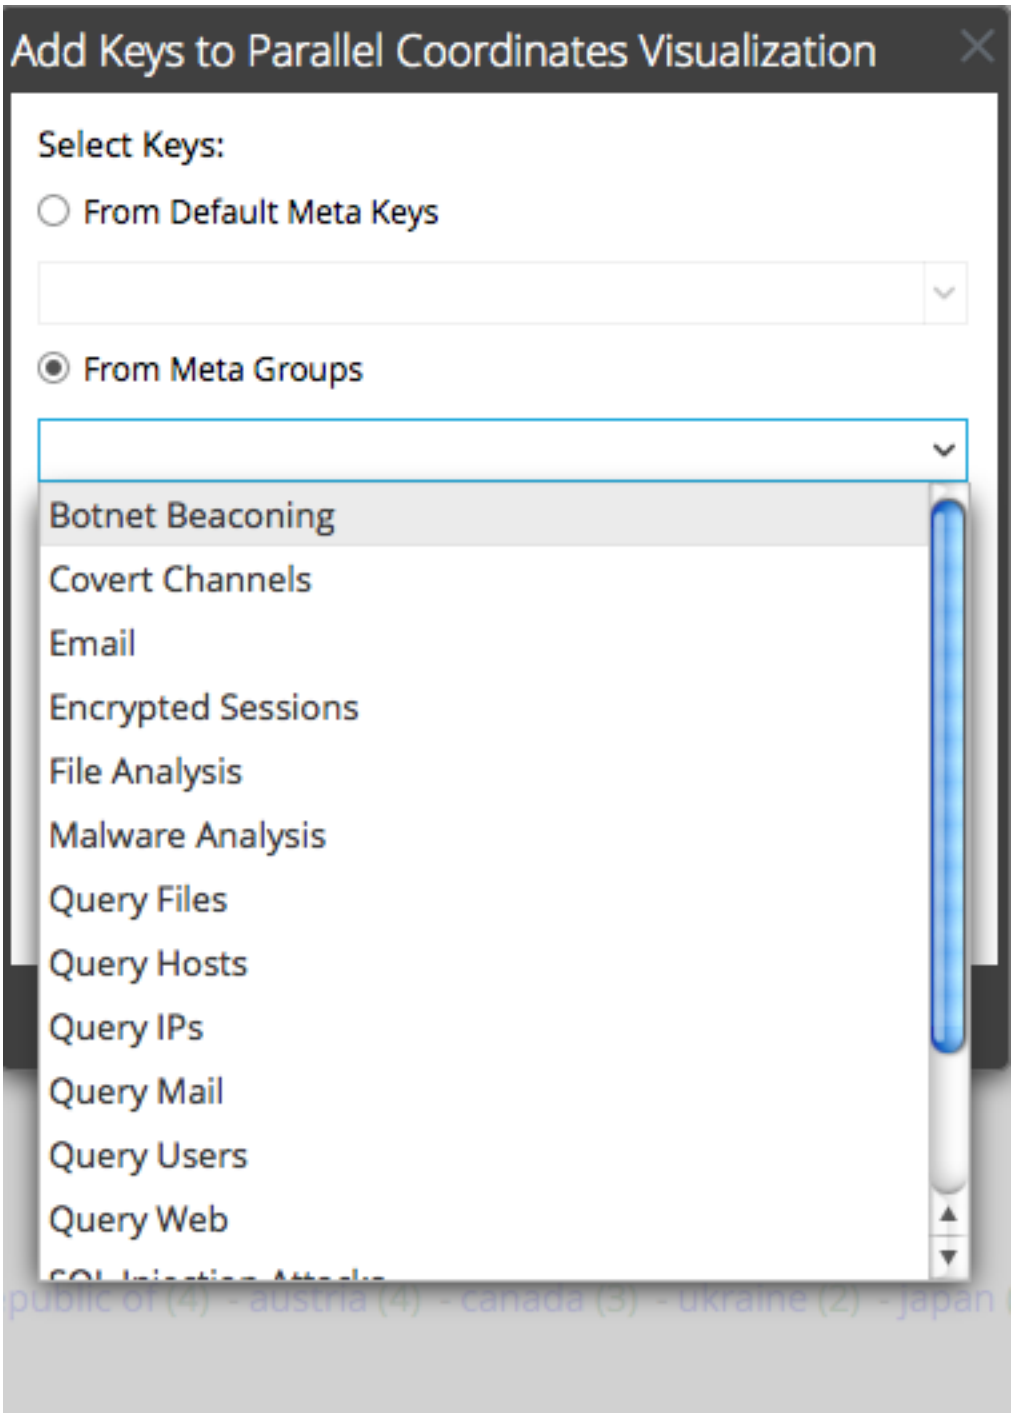
2. If you want to change the sequence of the meta keys, drag meta keys up or down to the desired sequence.
3. If you want to delete any meta keys, click in the selection box, and click .
   The meta keys are removed, but the change has not been applied.
4. If you want to revert to the previous state, click .
   Any meta keys you have deleted are restored and any changes that you made are removed.
5. If you want to select individual meta keys, click , select **From Default keys**, and in the drop-down list, select the meta keys.



The selected keys are listed.

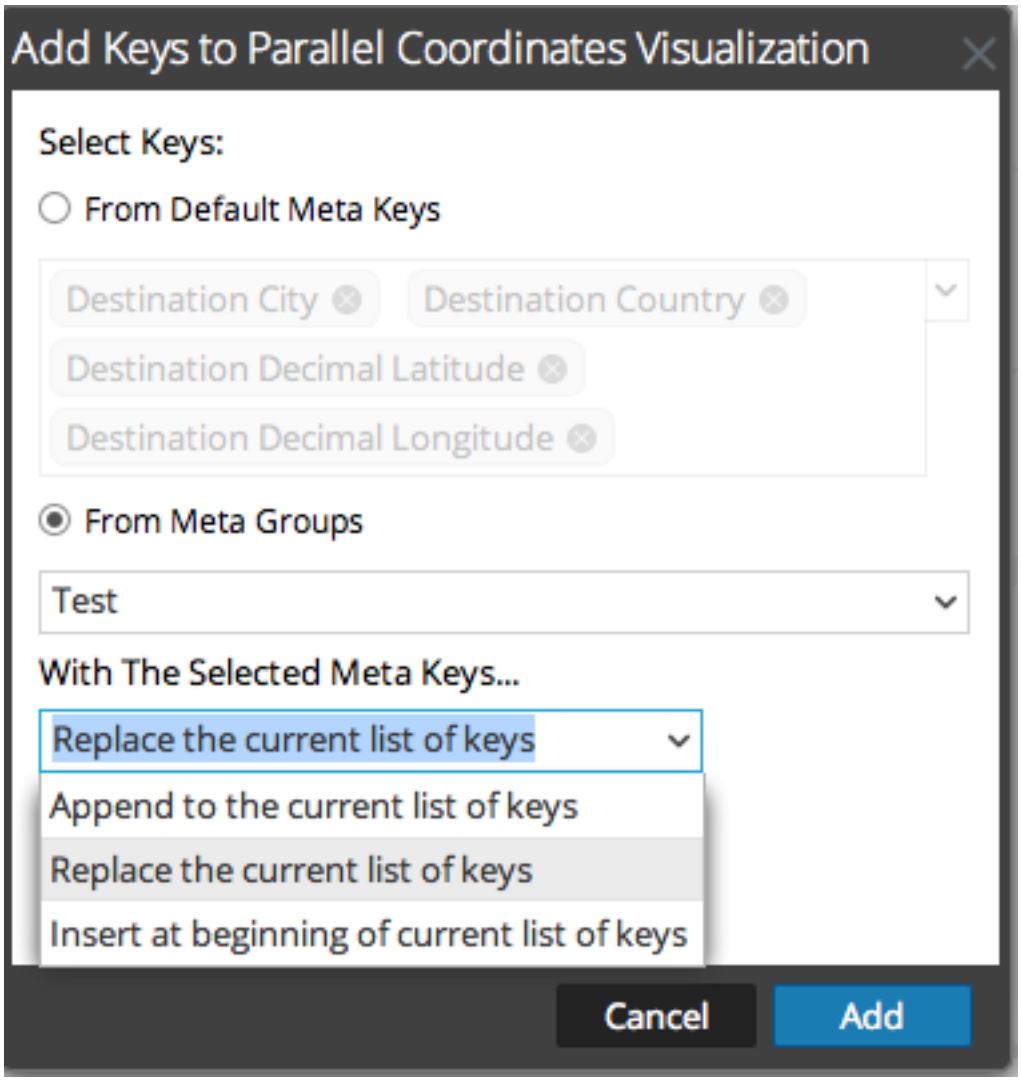


6. If you want to add all the keys in a meta group, you cannot add individual meta keys. Select **From Meta Groups**, and select a group from the drop-down list.



The selected meta groups are listed in the field.

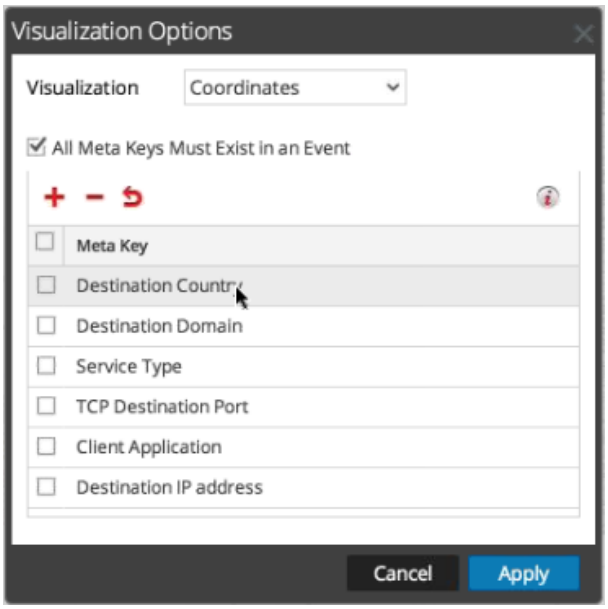
7. Select the method of adding the keys or groups: **Replace the current list of keys**, **Append to the current list of keys** (at the end), or **Insert at the beginning of current list of keys.**



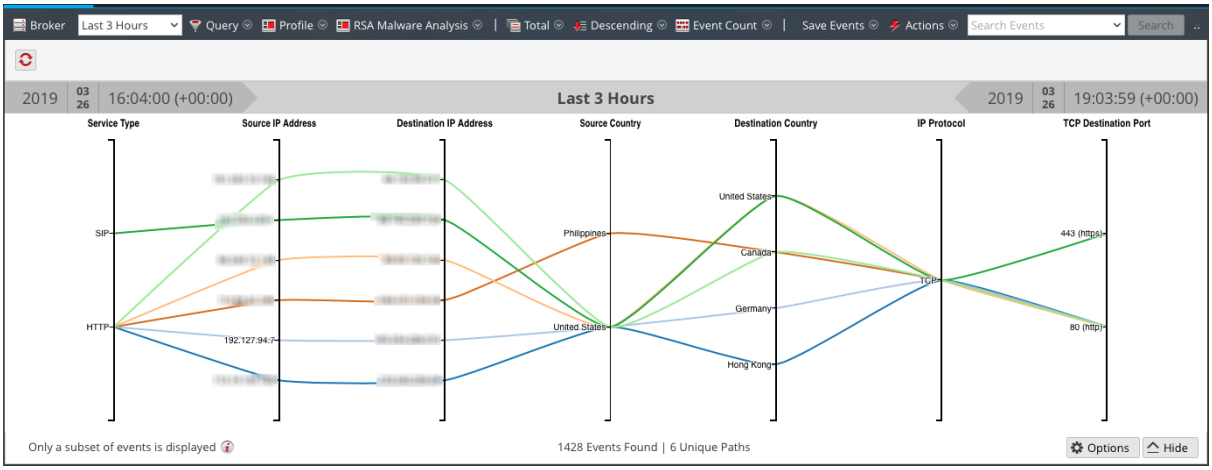
8. To complete the procedure, click **Add**.
   The Visualization Options dialog is displayed with the meta keys or groups you selected.
9. To display the new visualization chart, click **Apply**

Optimize a Parallel Coordinates Visualization

1. To optimize the visualization by removing events in which not all meta keys exist, select **Options.**

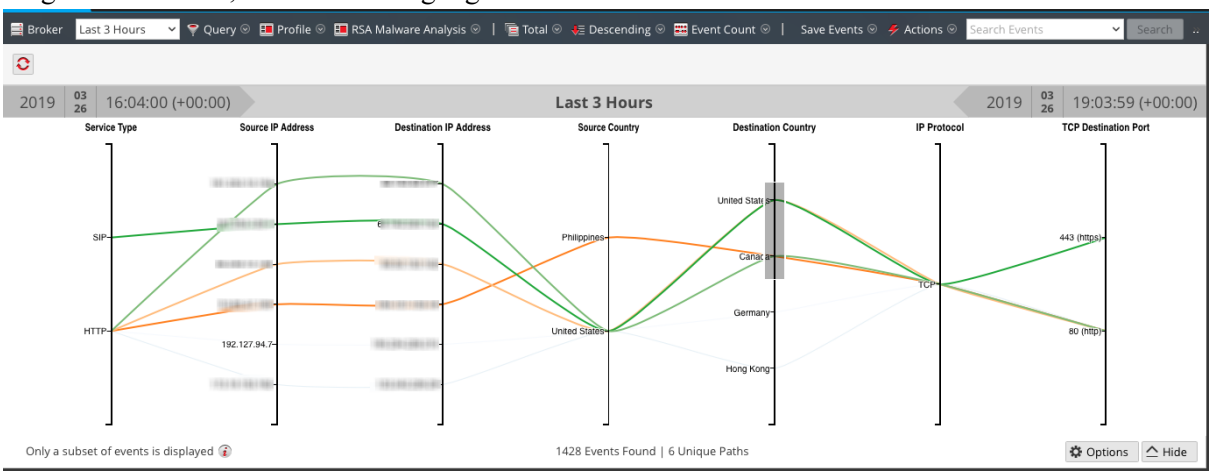


2. In the Visualization Options dialog, select **All Meta Keys Must Exist in an Event**. Click **Apply**. The resulting graph is more readable and useful and has fewer unique paths.



3. If you want to highlight a small set of points to see the path of the line from right to left, click on an axis. The cursor changes to cross hairs, which you can drag to select one or more values. When you

let go of the mouse, the lines are highlighted.



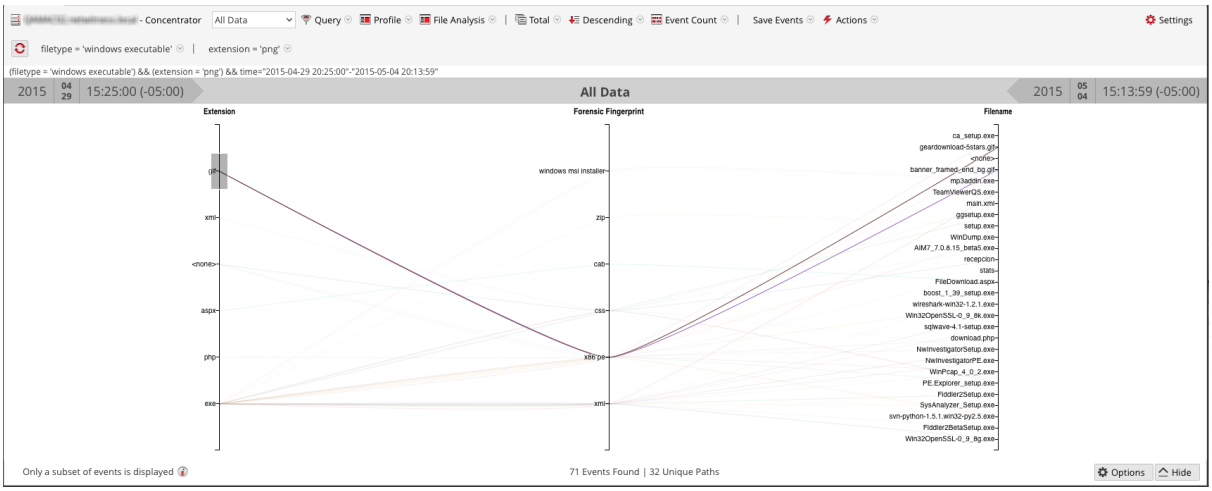
4. If you want to enlarge the visualization, drag the bottom edge of the panel down and drag the right edge of the browser window wider.

Sample Use Case

Below is an example of a parallel coordinates visualization of meta keys representing file metadata in a session. There are three meta keys or axes from left to right: Extensions, Forensic Fingerprint, and Filename with values listed along each axis. Values on the Extension axis show the file extension, and values on the Forensic fingerprint axis are windows executables. Normally the file type matches the expected forensics fingerprint; however, it is abnormal for a `gif` file type to be combined with the Windows executable fingerprint. The `gif` file type is selected to highlight the correlations of that file type, `x86pe` , and two filenames in the third axis so that an analyst can quickly identify the files that merit investigation.

To reach this view:

1. Order by Value and Sort in Ascending order.
2. Apply two filters (file type = 'windows executable' and extension = 'gif') in the Navigate view to limit the amount of data.
3. Configure a parallel coordinates chart by choosing three axes: `file extension`, `forensic fingerprint`, and `filename`.

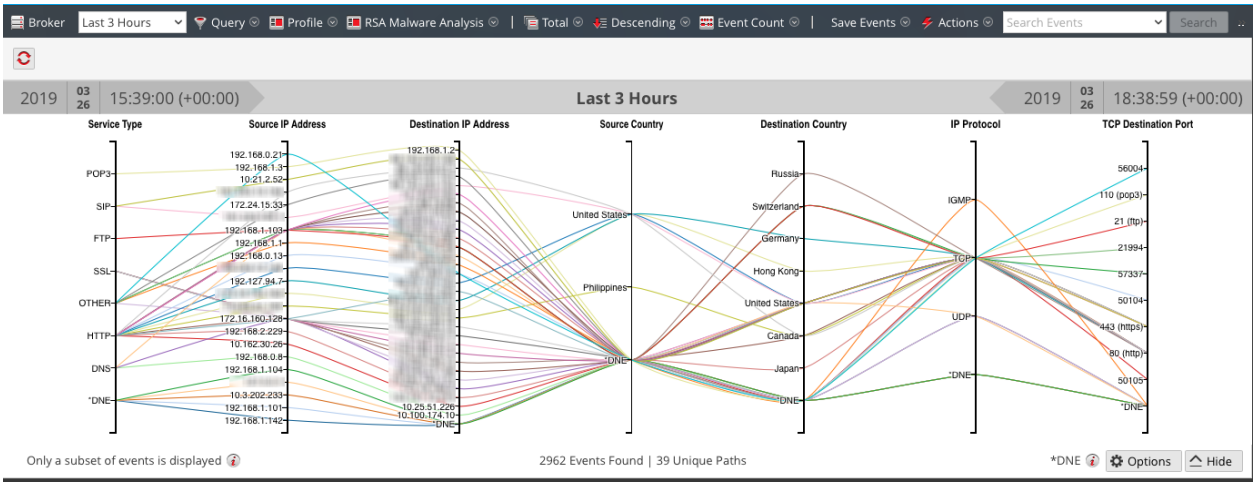


Sample Visualization of a Large Data Set

This example of a parallel coordinates visualization applied to a larger set of data illustrates several messages that help analysts to understand what has been charted.

- To create a chart, NetWitness begins scanning meta values and returning results. A typical time range could have up to 10,000,000 meta values. When the number of meta values returned reaches the Meta Values Result Limit, the chart is rendered even if NetWitness has not scanned a number of meta values equal to the Meta Values Scan Limit.
- There is a fixed limit on the amount of data that can be rendered as a parallel coordinates chart.The administrator configures parallel coordinates limits as part of the Investigation settings In the ADMIN > System view.

With a larger set of data, the parallel coordinates chart takes longer to process than the smaller set of data and meta keys. To preserve performance, NetWitness renders the meta values from the Values panel below until the limits set by the Administrator are reached. An informational message tells you: **Only a subset of events is displayed.**



Of all the data visualized for 2962 events, there were only 39 unique parallel coordinates paths. Some events are included though they do not include some of the meta keys; these are labeled **DNE** because the metadata does not exist in the event.

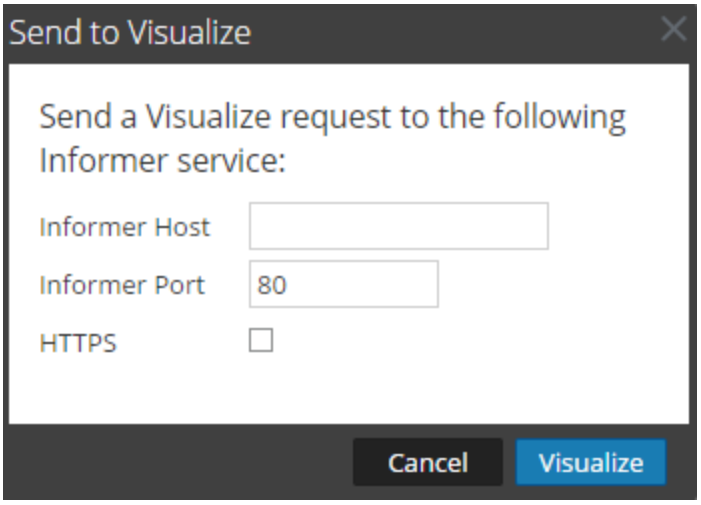
Visualize the Current Drill Point in Informer

This topic provides instructions for sending a drill point in the Navigate view to an Informer visualization.

Informer must be installed in your network and accessible by the service being investigated. You need to supply the host name and the port used on the Informer host to communicate with NetWitness.

To display a visualization in Informer of the current drill point:

1. With a drill point open in the Navigate view, click **Actions > Visualize**.
   The Send to Visualize dialog is displayed.



2. Type the Informer hostname or IP address, and verify the NetWitness server port used to communicate with the Informer host.
3. (Optional) Select the HTTPS option if the Informer host uses secured communications.
4. Click **Visualize**.
   The visualization is displayed in a new tab.

Downloading and Acting Upon Results

When working in Investigate, you may want to extract and shared data with other analysts, incident responders, SOC managers, and others. The topics in this section provide instructions for downloading results and creating incidents that appear in the Respond view:

- Download Data in the Events View
- Export or Print a Drill Point in the Navigate View
- Export Events in the Legacy Events View
- Add Events to an Incident in the Events View
- Add Events to an Incident in the Legacy Events View

Download Data in the Events View

In the Events view, you can download data from the Events panel and from a reconstruction. The Events panel download available in Version 11.4 and later is a bulk download of log and network events for all event types.

- Version 11.4.1 has the added capability to download visible metadata for all event types. From within a reconstruction, you can download events, logs, and files.
- Version 11.5 has the added capability to download metadata for all event types in the Events panel and in the event reconstructions.

Note: The information that you can view and download is managed by Role-Based Access Controls (RBAC) that your administrator has implemented. When RBAC is configured to prevent downloads of certain data, events for which you do not have download permission may appear to download successfully, but they are 0 byte in size. When RBAC is configured to prevent reconstruction of certain events, the reconstruction is disabled from the Events panel, but the bulk download button is still enabled.

Download Events or Metadata in the Events Panel

After submitting a query, you can download logs, network events, visible metadata (Version 11.4.1), or all metadata (Version 11.5) for events in your preferred format, directly from the Events panel. The preferences are set in the Event Preferences dialog and any changes made there are reflected inside the **Download** menu. See Configure the Events View for more information about preferences.

In the Events panel, you can select events individually or select all events returned by your search. The selection checkboxes appear only if you have permission to download the events. All checkboxes are deselected when a new query is submitted. When you select events and click Download, the Download menu is displayed. The number of events selected for each event type is displayed next to each option in the format `Events of this type selected/ Total number of events selected`. If an event type has no events selected, the corresponding download option is disabled and the number of events selected is displayed as `0 / Total number of events selected` as shown in the following figures. Version 11.5 has options to download all metadata or visible metadata, Version 11.4.1 menu has options to download visible metadata, and the Version 11.4 menu does not have those options.

If the Select All checkbox is selected in the Events list, the Download All options are available. In addition to the options that download all logs or network events, Version 11.5 has options to download all metadata or visible metadata. Version 11.4.1 menu has options to download visible metadata, and the Version 11.4 menu does not have the metadata options.

The difference between the All Meta options and the Visible Meta options are as follows:

- In Version 11.4.1 and later, visible metadata for the selected events is downloaded in the format that you selected in the Events Preference menu (**Visible Meta as Text**, **Visible Meta as CSV**, **Visible Meta as JSON**, or **Visible Meta as TSV**) or the format that you select under **Other Options** in the

Download menu at the time of download. The downloaded metadata for each event corresponds to the columns visible when the metadata is downloaded. The visible columns are determined by the selected column group and the Column Selector. For additional information about selecting columns, see Use Columns and Column Groups in the Events List. If the Summary column group is selected in the Events panel, all metadata for the events is downloaded. When you use one of the Download Visible Meta options; the downloaded metadata is sorted in order of collection time rather than the current sort order in the Events panel.

- In Version 11.5, all metadata for the selected events is downloaded in the default format that you selected in the Events Preference menu (**All Meta as Text**, **All Meta as CSV**, **All Meta as JSON**, or **All Meta as TSV**) or the format that you select under **Other Options** in the **Download All** menu at the time of download. The resulting download includes all metadata for the events selected, regardless of what columns are visible in the Events list. For example, if an event has 40 meta keys in the meta database, even if the column group in the Events list has 20 columns with 10 columns visible, all 40 meta keys for that event are included in the downloaded file.
- In Version 11.5.1, all metadata for the selected events is downloaded in the default format that you selected in the Events Preference menu (**All Meta as Text**, **All Meta as CSV**, **All Meta as JSON**, **All Meta as TSV**, or **Download Files**) or the format that you select under **Other Options** in the **Download All** menu at the time of download. The resulting download includes all metadata for the events selected, regardless of what columns are visible in the Events list. For example, if an event has 40 meta keys in the meta database, even if the column group in the Events list has 20 columns with 10 columns visible, all 40 meta keys for that event are included in the downloaded file.

Note: When you select all events for download, only the events in the current result set are downloaded. If you canceled the query before all results were returned, only the events that were loaded are downloaded.

To download event data for one event, multiple events, or all events in the Events panel:

1. Do one of the following:
   a. To select events individually, select the checkbox next to each event you want to download, and click the downward arrow on the **Download** menu button to see the options.

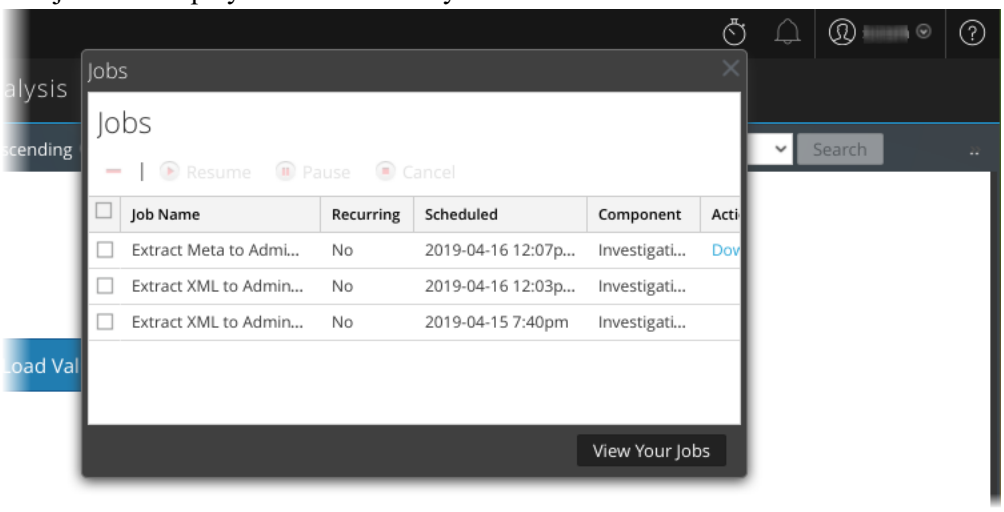
   b. To select all events displayed in the Events panel, select the checkbox at the top of the Events panel and click the **Download All** menu button.

2. Review the **Default Options** in effect in the top section of the menu. If you do not want to use the default format, you can choose a different format from the **Other Options** section of the menu.
   - Logs are downloaded in the preferred format that you selected in the Events Preference menu (**Logs asText**, **Logs as CSV**, **Logs as JSON**, or **Logs as XML**). If you want to choose a different format for this download, select one of the formats from **Other Options**.
   - Network events are downloaded as a PCAP. When downloading multiple network events in the Events panel, the format is always PCAP. The preferred format that you specified in the Events Preference menu (**Network as PCAP**, **Network as Payloads**, **Network as Request Payload**, or

Network as Response Payload) is ignored in this menu. Your preferred format applies only to downloading a single network event in the network reconstruction panel.

- Visible metadata is downloaded in the format that you selected in the Events Preference menu (**Visible Meta as Text**, **Visible Meta as CSV**, **Visible Meta as JSON**, or **Visible Meta as TSV**). If you want to choose a different format for this download, select one of the formats from **Other Options**. The downloaded metadata for each event corresponds to the columns visible when the metadata is downloaded. If the Summary column group is selected in the Events panel, all metadata for the events is downloaded.
- All metadata is downloaded in the format that you selected in the Events Preference menu (**All Meta as Text**, **All Meta as CSV**, **All Meta as JSON**, **All Meta as TSV**, or **Download Files**). If you want to choose a different format for this download, select one of the formats from **Other Options**. The downloaded metadata for each event includes all metadata, not just the visible columns.

3. Click the menu label: **Download** or **Download All**.
   The download begins immediately within the browser window if the **Download extracted files automatically** preference is set (**Events** view >). If the preference is not set, the download job for the selected events is added to the Jobs tray, where you can download the events.
   If the download fails, a message provides feedback regarding why the download failed. The download button is re-enabled and any selected events remain selected. These are examples of reasons for a failed download: timeout after X minutes, connection failed, event limit reached, and permission denied.
4. To display the Jobs tray, go to **Investigate > Navigate** or **Investigate > Legacy Events**, and click the  Jobs icon, which looks like a stop watch.
   The jobs are displayed in the Jobs tray.



Download a Log in the Text Reconstruction

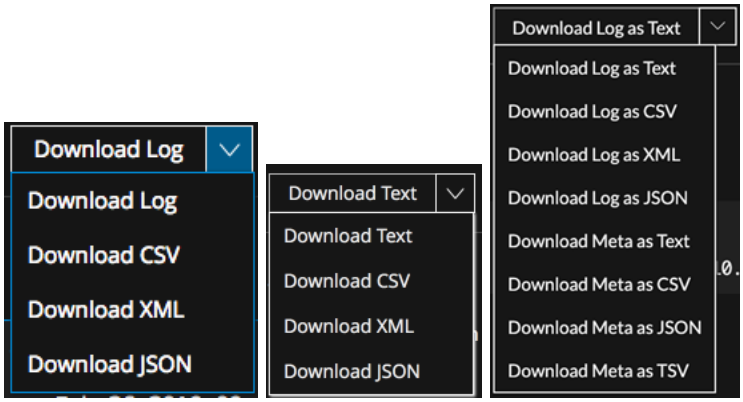
When viewing a text reconstruction of a log event, you can download a log file in the following formats using options in the Download Log menu:

- Raw log (log) using the **Download Log** (11.3) or **Download Text** (11.4 and later) or **Download Log as Text** (11.5 and later) option.
- Comma-separated values (CSV) using the **Download CSV** or **Download Log as CSV** (11.5 and later) option.
- Extensible Markup Language (XML) using the **Download XML** or **Download Log as XML** (11.5 and later) option.
- JavaScript Object Notation (JSON) using the **Download JSON** of **Download Log as JSON** (11.5 and later) option.

In Version 11.5 and later, you can also download the metadata for the log using one of these options:

Download Meta as Text, **Download Meta as CSV**, **Download Meta as JSON**, or **Download Meta as TSV**.

This is an example of the menu with the **Download Log** (11.3) or **Download Text** (11.4 and later) or **Download Log as Text** (11.5 and later) option. Version 11.5 has additional options to download metadata for the selected logs.



Note: For endpoint events, the **Download Log**, **Download Text**, or **Download Log as Text** option applies only to events that have at least one meta value exceeding 256 characters. For an endpoint event, the raw log is populated only when the meta value exceeds 256 characters. Long-running or historically downloaded files are not downloadable. For example, meta values like launch arguments can exceed 256 characters. In this case, 256 characters are available as a meta value while the full value is available in the raw log to view.

The downloaded log file contains the log and is named to help identify the service on which the log was collected, the session ID, and the file type. This is an example of the filename for a raw log: **Concentrator_SID2.log**. The exported log file is named using the following convention:

```
<service-ID or host name>_SID<n>.<filetype>
```

where:

- `<service-ID or host name>` is the name of the service (for example a Concentrator or Broker) where the session was saved.
- `SID<n>` is the session ID number.
- `<filetype>` identifies the format of the downloaded log. These are the possible log types: raw log, CSV, XML, and JSON. By default, the format is a raw log.

Note: Some formats do not have time stamps or the device IP where the event was generated, so a log downloaded in CSV, XML, or JSON format has an extra value called `timestamp` along with the raw log content. The additional information inside the log is in this form: `Log timestamp="1490824512" source="10.12.35.65".`

To download the log or the metadata for a log:

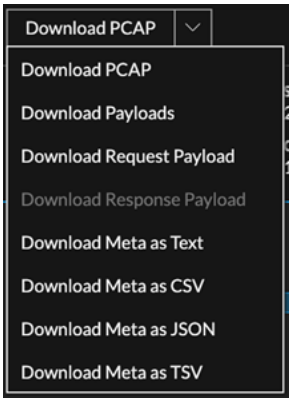
In the text reconstruction of a log event, do one of the following:

1. To download the log as a raw log (the default format), click **Download Log**, **(Download Text)** or **(Download Log as Text)**.
2. To download the log in one of the other formats, click the downward arrow on the **Download Log**, **Download Text**, or **Download Log as Text** button, then select one of the file formats for the downloaded log.
3. To download the metadata for the log, click the downward arrow on the **Download Log**, **Download Text**, or **Download Log as Text** label, then select **Download Meta as Text**, **Download Meta as CSV**, **Download Meta as JSON**, or **Download Meta as TSV**).

The log file or the metadata for the log is downloaded to your local file system in the format specified. If you initiate a download and move away from the view while the log is being extracted and before the log starts to download, the log is not downloaded in your browser. A message notifies you that you can find the downloaded log in the job queue.

Download Network Event Data in the Text or Packet Reconstruction

When viewing a packet reconstruction or a text reconstruction of a network event, you can export network data files for further analysis. In Version 11.5 and later, you can also download metadata for the reconstructed event.



The download includes events for the current time range and drill point. You can download the data in these formats:

- The entire event as a packet capture (`*.pcap`) file using the **Download PCAP** option.
- The payload as a `*.payload` file using the **Download All Payloads** (11.3) or **Download Payloads** (11.4) option.
- The request payload as a `*.payload1` file using the **Download Request Payload** option.
- The response payload as a `*.payload2` file using the **Download Response Payload** option.
- (Version 11.5) The metadata for the event using one of these options: **Download Meta as Text**, **Download Meta as CSV**, **Download Meta as JSON**, or **Download Meta as TSV**.

The label on the download menu button is one of these formats, based on the setting selected in the Event Preferences dialog. If the event does not have that type of data, the menu button is dimmed. You can click the downward arrow on the menu button to see which options are available. For example, if an event has a request payload, but no response payload, the Download Response Payload label is dimmed. You can click the downward arrow on the button and select **Download Request Payload** for this download. After selecting a valid format, clicking the button executes the download.

This is an example of the filename for a PCAP file: `C01 - Concentrator_SID1697309.pcap`. The exported network data file is named using the following convention:

`<service-ID or host name>_SID<n>.<filetype>`

where:

- `<service-ID or host name>` is the name of the service (for example a Concentrator or Broker) where the session was saved.
- `SID<n>` is the session ID number.
- `<filetype>` is pcap, payload, payload1, or payload2.

The network data is downloaded directly into your browser if the download is quick. If the download takes longer due to network factors or file size, the file is downloaded in the background and the task is tracked in the Jobs queue. In this case, you can check your jobs in the queue and get the file when the download is complete.

Note: If you initiate a download and move away from the view while the file is being extracted and before the file starts to download, the file is not downloaded in your browser. A message notifies you that you can find the downloaded document in the job queue.

To export an event as a network data file or to download the metadata for the event:

Go to the packet reconstruction of a network event and do one of the following:

1. To download the event as a PCAP file (the system-defined default format) or in the user-defined default format, click the **Download <format>** button. The label is the same as the download option set in the Events Preferences dialog.
2. To download the event in one of the other formats, click the downward arrow on the button, and select one of the file formats for the downloaded event data.

3. To download the metadata for the event, click the downward arrow on the button, and select one of the file formats for the downloaded metadata.

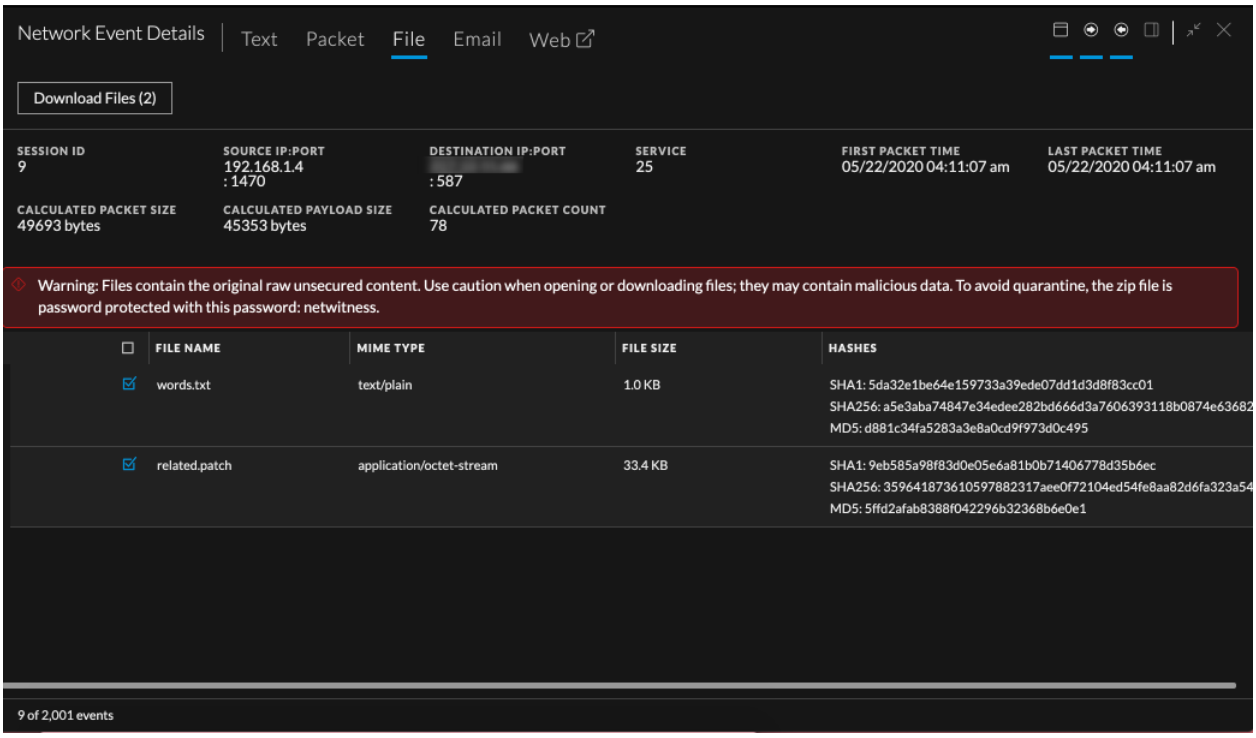
The network data file is downloaded to your local file system in the format specified or the metadata for the event is downloaded in the format specified.

Download Files from a Network Event in the File Reconstruction

When viewing reconstructed network events that contain files in the file reconstruction, you can select one or more files, or all files, to download to your local file system.

Note: If you initiate a download and move away from the view while the file is being extracted and before the file starts to download, the file is not downloaded in your browser. A message notifies you that you can find the downloaded file in the job queue.

When files are selected, the Download Files button becomes active and reflects the number of files selected.



Clicking Download Files exports the selected files as a password-protected zip archive. The password to open the exported archive is `netwitness`. Exporting the files in this form ensures that:

- The archive is not quarantined by antivirus software.
- Potentially malicious files are not automatically opened by the default application and executed.

When downloading files form the file reconstruction, the exported archive is of the form `<service-name>SID<service ID><file-count>_FILES_FILES`, for example, `Broker_SID8_1_FILES_FILES.zip`. This is the password to open the zip archive: netwitness.

`<service-ID or host name>_SID<n>_<file-count>FILES_FILES.zip`

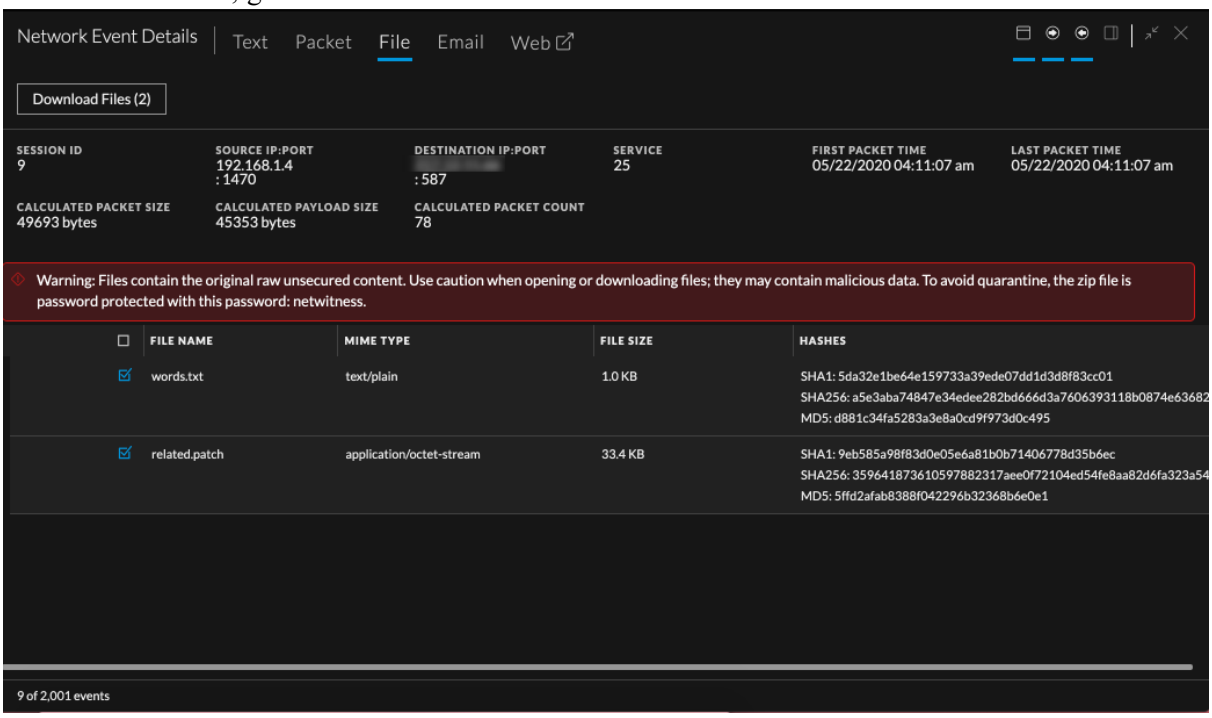
where:

- `<service-ID or host name>` is the name of the service (for example a Concentrator or Broker) where the session was saved.
- `SID<n>` is the session ID number.
- `<file-count> FILES` is the number of files in the archive.
- `FILES` identifies the reconstruction type from which the files were downloaded.

Caution: Caution is advised when unzipping and opening files that are associated with a default application; for example, an Excel spreadsheet may automatically open in Excel before you have a chance to verify it is safe.

To export files in a reconstructed event:

1. In the **Events** view, go to the file reconsruction of an event that contains files.



2. Click one or more files that you want to extract, and click **Download File** or **Download Files**. The job is scheduled and when complete the selected file are downloaded, in the form of a password-protected zip archive, to the local file system.
3. To open the archive on your local file system, enter the following password when prompted: `netwitness`.

Download Attachments from an Email Reconstruction

When viewing an email reconstruction that contains attachments, you can select one or more attachments, or all attachments (Version 11.4.1.x), to download to your local file system. This feature exports the selected files as a password-protected zip archive. The password to open the exported archive is `netwitness`. Exporting the files in this form ensures that:

- The archive is not quarantined by antivirus software.
- Potentially malicious files are not automatically opened by the default application and executed.

When downloading files from an email reconstruction, the filename is of the form `<service-name>_SID<n>_EMAIL`, for example, `Broker-_SID34_EMAIL.zip`. This is the password to open the zip archive: netwitness. The exported archive is named using the following convention:

`<service-ID or host name>_SID<n>_EMAIL.zip`

where:

- `<service-ID or host name>` is the name of the service (for example a Concentrator or Broker) where the session was saved.
- `SID<n>` is the session ID number.
- `EMAIL` is the type of reconstruction from which the files were downloaded.

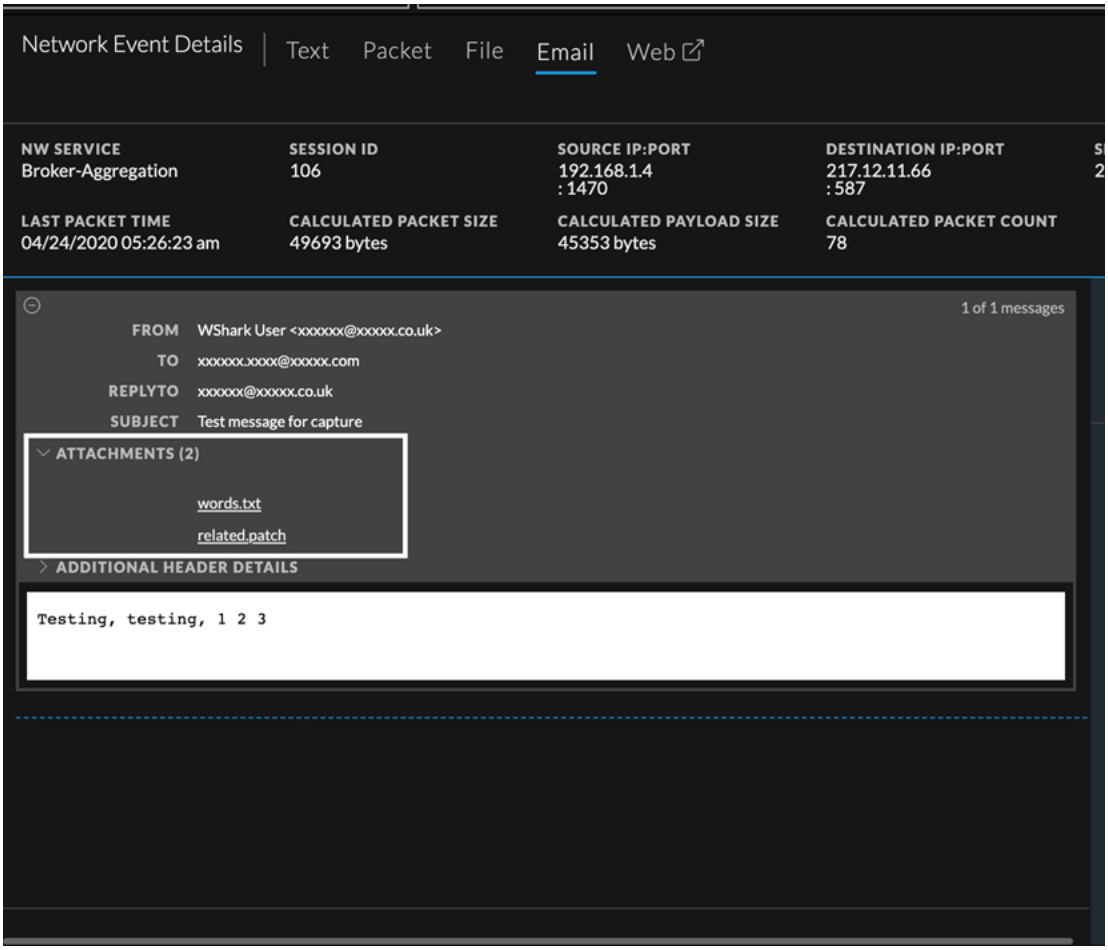
Caution: Caution is advised when unzipping and opening files that are associated with a default application; for example, an Excel spreadsheet may automatically open in Excel before you have a chance to verify it is safe.

To download email attachments:

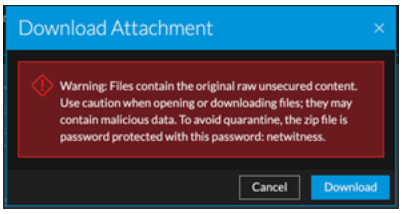
1. Go to the **Events** view and open an email reconstruction by clicking an event that contains an email with attachments.

2. Expand the **Attachments** drop-down list and do one of the following:

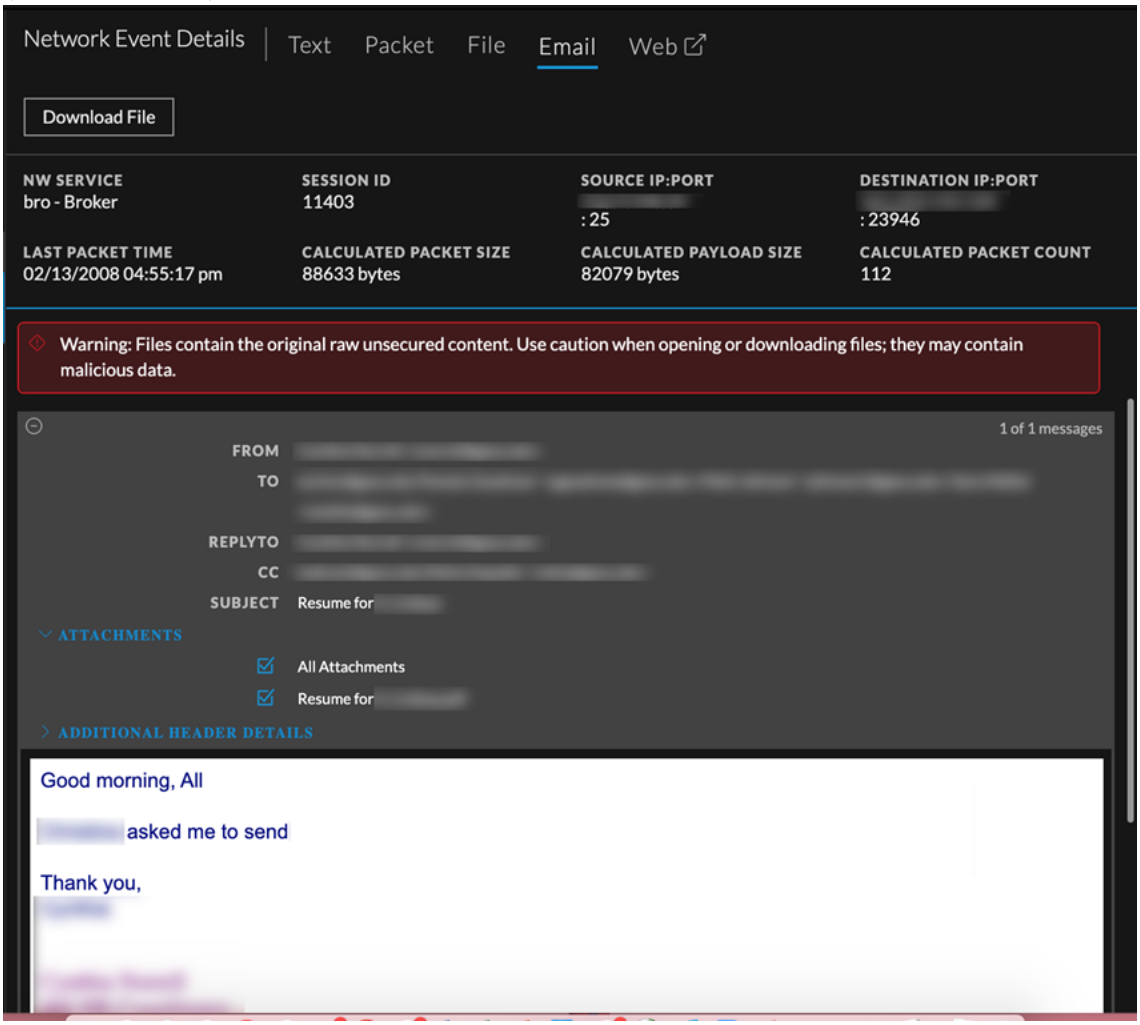
   a. (Version 11.5 and later) Click a link to an attachment.



A dialog warns you that downloaded email attachments may contain malicious data and asks you to cancel or confirm the download. If you want to complete the download, click **Download**. Otherwise, click **Cancel** to cancel the download.



b. (Version 11.4.1.x) Select one or more attachments or **All Attachments**.



A warning message is displayed in the reconstruction. Click the **Download File** or **Download Files** button. The attachments are downloaded with no additional opportunity to cancel.

Export or Print a Drill Point in the Navigate View

In NetWitness Investigate, when the data for a drill point is displayed in the Navigate view, you can:

- Extract files from a session and choose the type of files to extract: archives, audio BitTorrent, documents, executable, images, other, video, and web.
- Export the drill point as a packet capture (PCAP) file, a log file, or a metadata file.
- Print the drill point.

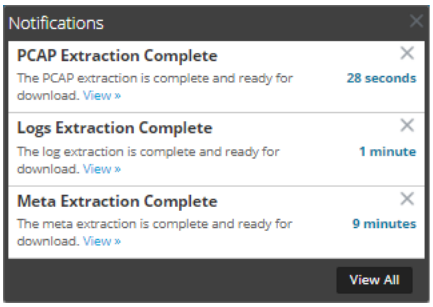
The details being exported are affected by both the time range and drill point at the time of exporting.

Note: When you export the drill point as a log file, only the log sessions are exported. The job queue message refers to the total number of sessions in the drill point rather than the number of logs. For example, if the drill point has 505 sessions and only five log sessions, the job queue message states that NetWitness is extracting logs for 505 sessions.

To export a drill point from the Navigate view:

1. Conduct an investigation until you reach the desired drill point.
2. For Version 11.0, In the toolbar, select **Actions > Export** and select one of the export options: **PCAP**, **Logs**, or **Meta**.
   The drill point is extracted, and a message advises that the job is scheduled. You can check the jobs page for the status.
3. For Version 11.1, in the toolbar, select **Save Events** > and select one of the export options: **PCAP**, **Logs**, **Files**, or **Meta**.
   A dialog gives you an opportunity to edit the default filename for the file. The default filename is in the form `investigation-Feb-21-15-44-33`. When you are exporting a PCAP, the file is exported with no choice of formats. If you are using one of the other export options, a dialog is displayed.
4. In the dialog, select:
   - The export log format: **Text**, **XML**, **CSV**, or **JSON**.
   - The file types to export: Archives, Audio, BitTorrent, Documents, Executables, Images, Other, Video, and Web.
   - The Meta format: **Text**, **CSV**, **TSV**, **JSON**.

5. When the scheduled file extraction is complete, it is displayed in the Job Notifications tray.



6. Click the **View** link in the Jobs tray and download the specific extraction file requested.

To print the current drill point:

In the Navigate view, you can display the contents of the current drill point in printer friendly format in the browser window.

To display the current drill point in a print view:

1. With a drill point open in the **Navigate** view, select **Actions** > **Print** in the toolbar.
   A new tab is created with the print view of the current drill point.



2. Use the print option in your browser to send the printable view to the printer.

Export Events in the Legacy Events View

In the Legacy Events view, the Actions menu has an option to export events from the event being viewed to an archive.

Note: You can only export files that you have permission to view or access.

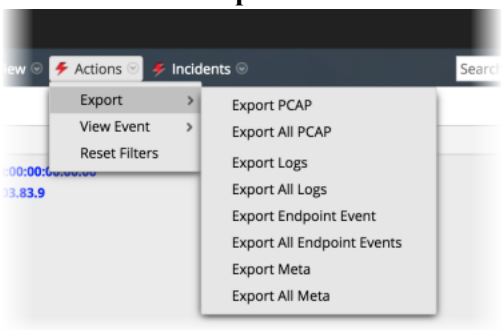
The export function queries the service for all sessions inside the selected time range and drill point to extract the content of each session. The details being exported are affected by both the time range and drill point at the time of exporting. In the File Extraction dialog, you can choose to export:

- PCAPs
- Logs
- NetWitness Endpoint events
- Meta values

The format of the exported archive: ZIP or GZIP file. After you send the request, a job is scheduled and you can track the job in in the Jobs tray. If there is an error retrieving the log or PCAP from the service, an error notification is displayed.

To extract files from an event:

1. While in the **Event view**, click an event.
2. Click **Actions > Export.**.



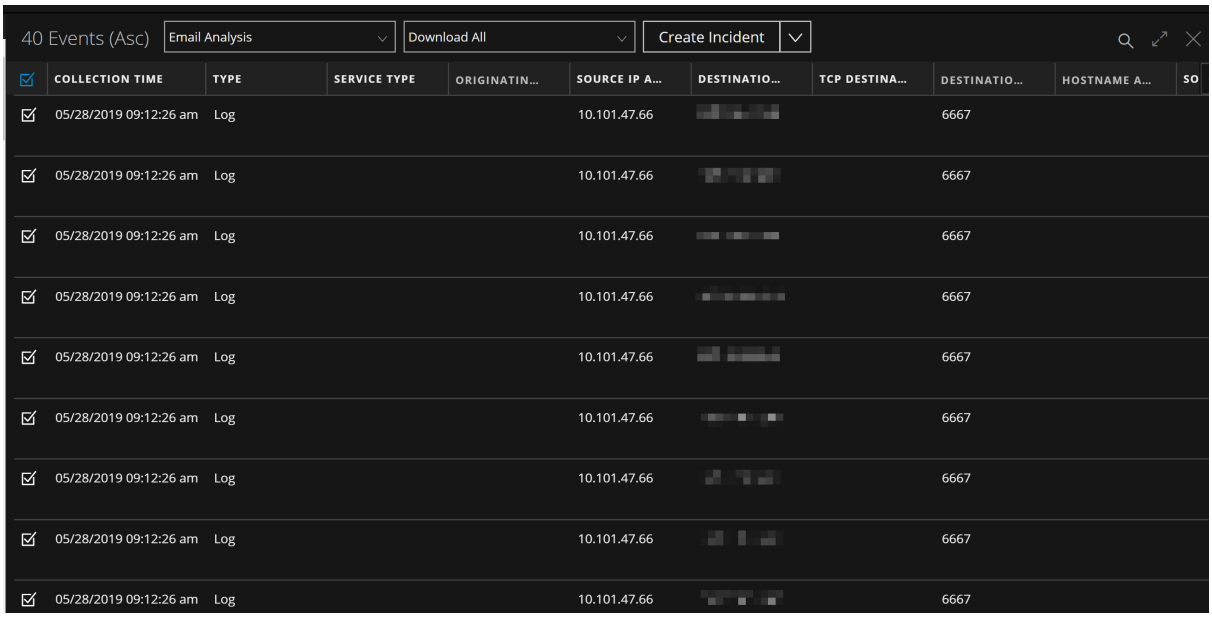
3. Select the export option.
   A message informs you that the PCAP is being downloaded.

Add Events to an Incident in the Events View

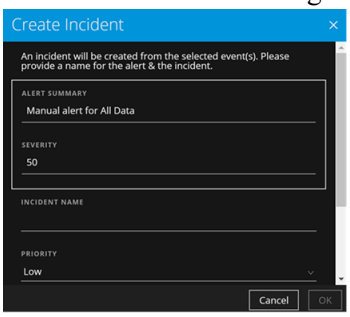
When conducting an investigation in the Events view, you can select one or more events and create an incident that is available for incident responders in Respond. When you create an incident, if access restrictions are in effect, you can view only incidents to which you have access. For example, when creating incidents from the Investigate view, analysts must assign the incidents to themselves to view them in the Respond view. You can also add events to an existing incident in Respond to which you have access.

Note: An administrator must configure the `respond-server.incident.manage` and `investigate-server.incident.manage` roles and permissions. For more information, see "Role Permissions" and "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*.

1. Go to **Investigate > Events**.
2. In the Events view, select one or more events.



3. Click **Create Incident.**
   The Create Incident dialog is displayed. Complete the information in the Create Incident dialog.



a. Select the severity. The alert summary field is a pre-defined value which is auto-populated but can be edited if required.

b. Type a name for the incident in the **Incident Name** field.

c. From the **Priority** drop-down list, select a priority for the incident. For example, an incident may be critical, high, medium, or low priority.

d. Select an assignee for the incident from the drop-down list. This list includes the built-in users that have access to Investigate as well as any custom users that have been added to your system. For example, this list might include users for admin, analyst, dpo, operator, and users for incident responders.

e. From the Categories drop-down list, select one or more categories of events that apply to this incident.

f. Click **OK**.
   An incident is created with the selected event in Investigate.

4. To add one or more events to an existing incident, select one or more events, and then click **Add to Incident**.

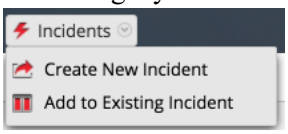
5. In the Add to Incident dialog, select the alert summary and severity, and select one or more open and existing incidents to which the incidents will be added. You can Search for an existing incident by Incident-ID or Incident Name. When ready, click **OK**. The event is added to the selected incidents and updated in Respond.

Add Events to an Incident in the Legacy Events View

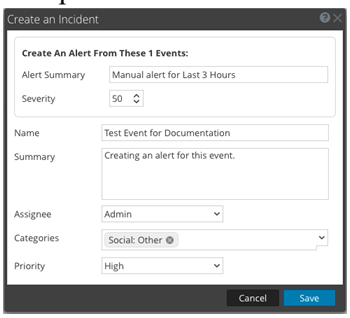
When conducting an investigation in the Legacy Events, you can select one or more events and create an incident that is available for incident responders in Respond. When you create an incident, if access restrictions are in effect, you can view only incidents to which you have access. For example, when creating incidents from the Investigate view, analysts must assign the incidents to themselves to view them in the Respond view. You can also add events to an existing incident in Respond to which you have access.

Note: An administrator must configure the required roles and permissions as described in "Role Permissions" and "Manage Users with Roles and Permissions" in the *System Security and User Management Guide*.

1. Go to **Investigate > Legacy Events**.
2. In the Legacy Events view, select one or more events, and then **Incidents > Create New Incident**.



3. Complete the information in the Create an Incident dialog.



   a. Select the severity, an integer between 1 and 100, with 100 being the most severe.
   b. Type a name for the incident and describe the incident in the **Summary** field.
   c. Select an assignee for the incident from the drop-down list. This list includes the built-in roles that have access to Respond as well as any custom roles that have been added to your system. For example, this list might include roles for admin, analyst, dpo, operator and roles for incident responsers.
   d. From the **Categories** drop-down list, select one of more categories of alerts that apply to this incident.
   e. From the **Priorities** drop-down list, select a category for the incident. For example, an incident may be critical, high, medium, or low priority.
   f. Click **Save**.
   The new incident is created and is available immediately in the incident queues for the selected role in Respond.

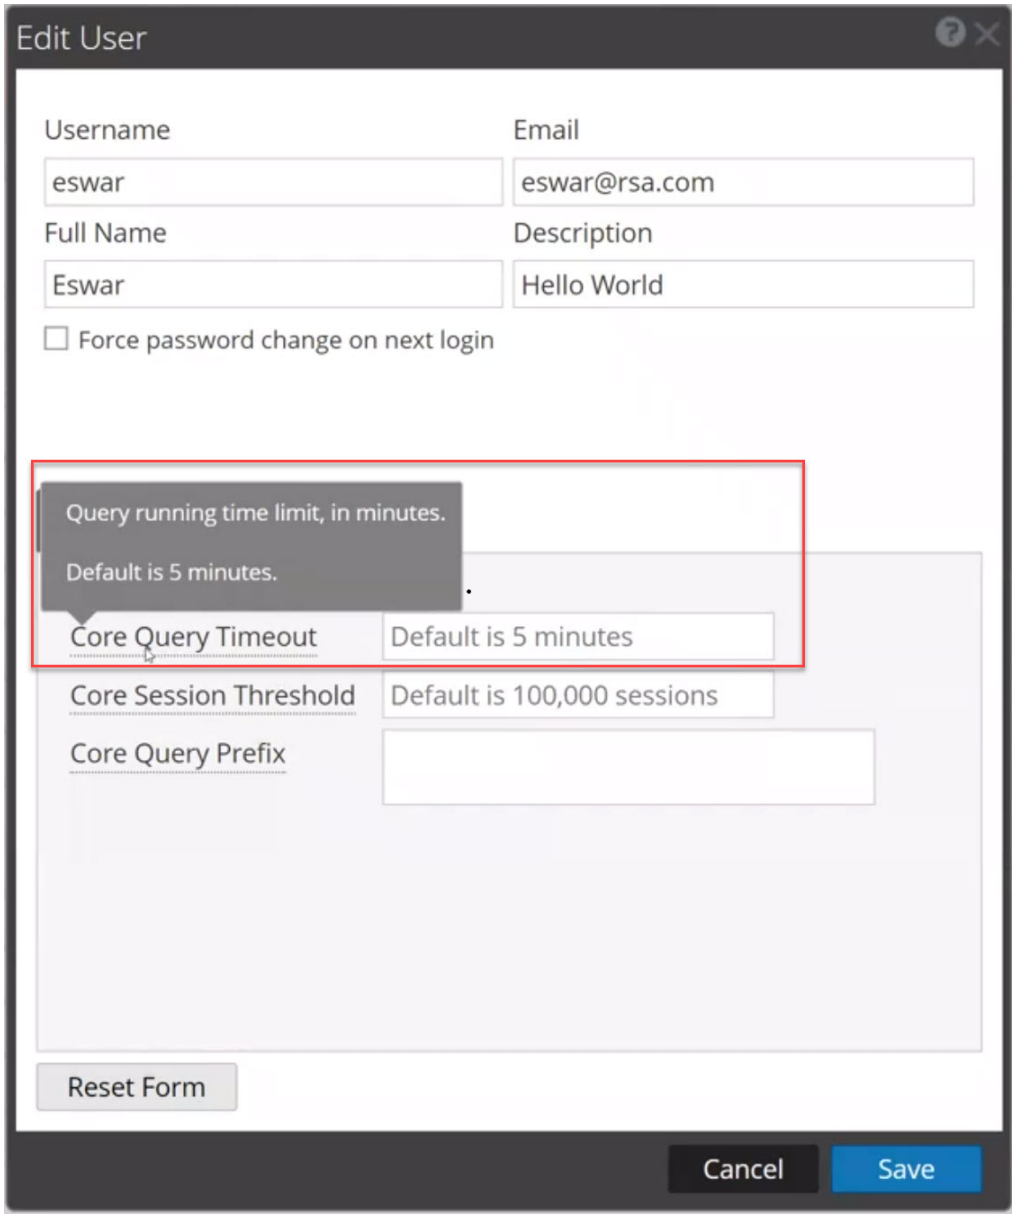
4. To add one or more events to an incident, select one or more events, and then **Incidents > Add to Existing Incident**.
5. In the Add Events to an Incident dialog, select the severity, and select one or more incidents to which the events will be added. You can Search for an existing incident by Incident-ID or Incident Name. When ready, click **Add to Incident**.
   The events are added to the selected incidents and updated in Respond.

Troubleshooting NetWitness Investigate

This section provides information about possible issues when using NetWitness Investigate.

PCAPS and SDK Issues

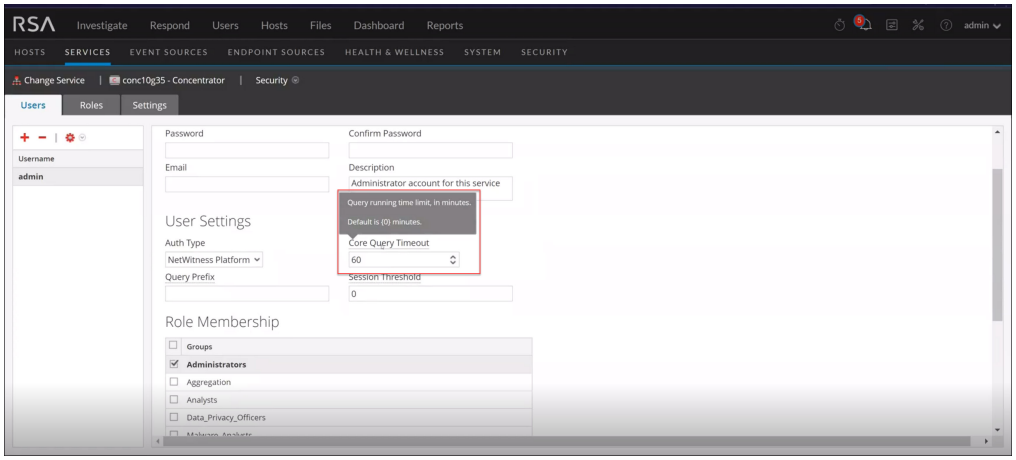
Behavior	When you extract a large set of pcaps or run SDK Queries with huge matching sessions, the query fails or the session times out.
Issue	All users have default timeout of 5 minutes, when this time exceeds, the query stops running. Also, the Core service account has a default timeout of 60 minutes.
Explanation	Verify the query time for the user and the service accounts. You can modify the default settings from the user interface based on your requirement: To edit the timeout settings for a user: 1. Log in to NetWitness Platform, 2. Navigate to **Admin** > **Security** 3. Select a user you want to edit. 4. Click the **Attributes** tab. 5. Edit the default **Core Query Timeout** settings.



6. Click **Save**.

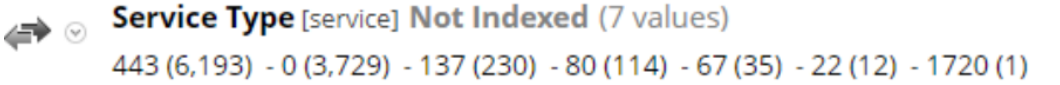
To edit the timeout settings for a core service account:

1. Log in to NetWitness Platform.
2. Navigate to **Admin** > **Security**.
3. Select a user you want to edit.
4. Click the **Attributes** tab.
5. Enter the credentials in the **Password** and **Confirm Password** fields.



6. Click **Save**.

Navigate View and Legacy Events View Issues

Behavior	A meta key that normally returns values in the Navigate view returns values, but has a `Not Indexed` message following the meta key name. For example, the in this figure the Service Type meta key is followed by the message: `Service Type[service] Not Indexed`. 
Issue	When you first set up the environment or very rarely after performing a data reset on the broker due to other issues, you see meta keys as Not Indexed when they are indexed at meta key or meta values level.
Explanation	To fix the issue on a Broker, log out of NetWitness Platform and then log in again. Valid sessions will be displayed.

Message	`Not indexed; will experience longer than usual load times.` in the Manage Meta Groups dialog.
Issue	Meta keys in the Manage Meta Groups dialog are marked by a red exclamation point, and the error message is displayed. This can occur when investigating a Broker or Decoder while adding a meta group with meta keys that are not indexed in the index file or the custom index file for the service. For a Broker, it could mean that the Broker has not begun aggregating data from a Concentrator. In this case the Broker will not have the contents of the custom index file from the aggregate services and the keys will not be indexed. For a Decoder, it means that the meta keys are not indexed in the Decoder index or custom index file.

Explanation	To fix the issue on a Broker, log out, log in, and restart the Broker service so that it can aggregate the meta key information from connected Concentrators. To fix the issue on a Decoder, edit the custom index file to index the meta keys, log out, log in, and restart the Decoder service.

Behavior	When downloaded from the Event Reconstruction view, logs and metadata are always in text format irrespective of the format selected in the Legacy Events view.
Issue	When you download metadata or a log in the Event Reconstruction view, the format that you selected in the Legacy Events view is not used. The exported data is always in text format.
Explanation	Download metadata and logs from the Legacy Events view if you want to use a format other than text format.

Events View Issues

Message	`Applicable for hosts with 4.x Endpoint agents installed, please install the NetWitness Endpoint Thick Client.`
Issue	When you click **Pivot to Endpoint** in the Events view, no data is displayed and the message is displayed.
Explanation	Version 4.4 of the NetWitness Endpoint Thick Client must be installed on the same server, the NWE meta keys must exist in the `table-map.xml` file on the Log Decoder, and the `index-concentrator-custom.xml` file on the Concentrator. The NWE Thick Client is a Windows only application. Complete setup instructions are provided in the *NetWitness Endpoint User Guide* for Version 4.4.

Behavior	Download jobs are in a Waiting state or Failed state in the Jobs tray during and after upgrading the software to Version 11.4.
Issue	If you had download jobs running while your administrator was upgrading the software, you may see a job in a Waiting state while the upgrade is in progress and then in a Failed state after the upgrade is complete. You cannot resume or cancel the failed job.
Explanation	To delete the failed jobs, select the failed jobs in the Jobs tray and click .

Message	Event counts in the Filter Events panel and the Events panel may sometimes differ when showing results for the same query.
Issue	The Filter Events panel uses only index data to produce counts of events, which is less

	accurate than the Events panel. The Events panel results are filtered for exact matches on data retrieved from the meta database, which takes longer to process.
Explanation	At worst the difference is in false positives in the Filter Events panel, not false negatives; so you will not miss an event.

Message	`Event Analysis requires all core services to be NetWitness 11.1. Connecting prior versions of services to the 11.1 NetWitness Server results in limited functionality (see "Investigate in Mixed Mode" in the *Physical Host Upgrade Guide*).`
Issue	When attempting to investigate a service that has not been updated to Version 11.1 in the Event Analysis view, the informational message is displayed.
Explanation	When an analyst opens the Event Analysis view in mixed mode (that is, some services are upgraded to 11.1 and later, and some are still on 11.0.0.x or 10.6.x), Role-Based Access (RBAC) is not applied uniformly. This affects viewing and downloading content, and validation of filters in the interactive breadcrumb. You will see this informational message when you open Events. As you select a service, services that are not up to date are displayed in a red box, with the message that the service is not up to date. When your administrator has upgraded all connected services to 11.1 and later, these features work as expected.

Message	`Forbidden. You cannot access the requested page.`
Issue	When attempting to access the Events view, the view opens with the message.
Explanation	Your administrator has prevented access to the Events view using role and permissions.

Behavior	If you can download an event in the Events view, but get a 0-byte file, the administrator may have restricted access to the content.
Issue	Role-Based Access Controls applied by your administrator allowed you to download an event for which you did not have permission; therefore, the file download was empty.
Explanation	If you believe you should have access to the event, contact your administrator.

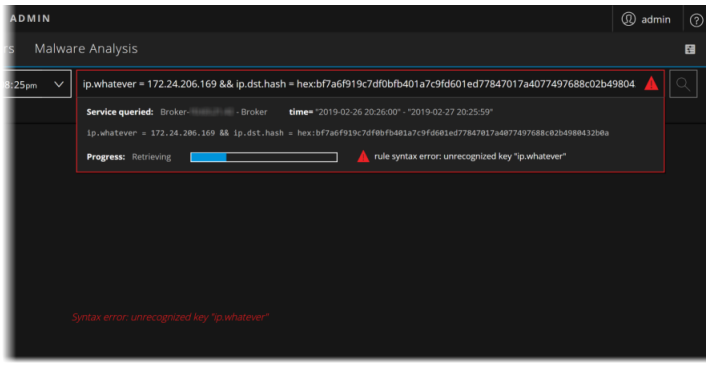
Message	`Insufficient permissions for the requested data.`
Issue	While attempting to access an event in the Events view, the message is displayed.
Explanation	You have entered an event ID for an event that you do not have permission to view. The administrator may have placed some restrictions to limit access by role and permissions.

Message	`Invalid session ID: <<eventId>>`
Issue	No `sessionId` matches the `sessionId` that you queried.
Explanation	The reason for an invalid session ID can vary. Perhaps you edited the session ID manually, and no such session exists. Another case may be when you query a Broker, and the aggregated data has not been refreshed, you may see this error for a session that no longer exists.

Behavior	Investigation Profiles and built-in column groups are not present in 11.1 Event Analysis.
Issue	Post upgrade to RSA NetWitness v11.1, the default column groups - Endpoint Analysis, Outbound SSL and Outbound HTTP are not added under column groups. Also, a few of the Investigation Profiles are missing post upgrade.
Explanation	It is observed that this issue occurs only when you have created a custom column group with the name which is same as one of the new 11.1 OOTB custom column group name. For example, if you create a custom column group in 11.0 with name **RSA Endpoint Analysis** then after upgrade to 11.1. Due to the same name already existing in 11.1, OOTB column groups and built-in profiles will not be available in the UI. To fix this, change the name of custom column group to something other than one of the OOTB column groups and restart the jetty server by using the following command on the NetWitness server: `systemctl restart jetty`

Message	`Memory limit of <XXXXXX> GB reached, controlled by setting max.query.memory`
Issue	The query that you submitted failed because the result set was too large, and the memory limit set by `max.query.memory` was reached.
Explanation	To avoid this error, try to further limit results by narrowing the time range, adding filters, and decreasing the number of columns in the column group. You can also ask an administrator to limit the number of events returned.

Behavior	No text data was generated during content reconstruction. This could mean that the event data was corrupt or invalid, or that an administrator has disabled the transmission of raw endpoint events in the Endpoint server configuration. Check the other reconstruction views.
Issue	When you reconstruct an event as text in the Events view, no data is displayed and the message is displayed.
Explanation	If you do not see the raw text in other Events views or Legacy Events view reconstructions, and you believe the data is not corrupted or invalid, your administrator has likely disabled transmission of raw endpoint events on the NetWitness Endpoint server. Contact your administrator for additional information.

Message	`Rule Syntax error: Unrecognized key "<meta key or meta entity name>"` `Syntax error: Unrecognized key "<meta key or meta entity name>"`
Issue	While querying a service, the matching events are not listed and the message is displayed in the query console and the Events view. 
Explanation	The query you entered is querying a meta entity that is not configured properly. All upstream devices connected to the Broker being queries should have the same entity configuration. This error indicates that the Broker is operating with mismatched entity definitions. Ask your administrator to review the configuration described in "Index Customization" in the *Core Database Tuning Guide*.

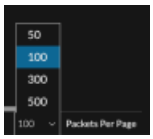
Message	`Selected Column Group is no longer available. The default summary column group has been selected instead.`
Issue	If you had set a preferred column group before the 11.4 upgrade, on your first visit to the Events view, the flash message is displayed even when the column group is available or is the default group (summary). This issue was resolved in Version 11.4.1.
Explanation	This is a one-time occurrence. If you reload the Events view, the message is not displayed.

Message	`Session is unavailable for viewing.`
Issue	While querying an event ID, the reconstruction is not displayed and the message is displayed.
Explanation	The query you entered is trying to look at restricted data, for example, if you are allowed to see only log data and you are using a link to network data .

Message	`The query on channel <channel-number> was auto-canceled by the system for exceeding time usage limits. Check timeout values. Query running time was 00:05:00 (HH:MM:SS)`
Issue	If you continually get this timeout message, first check the query console to determine if there are issues around time it takes for a service to respond, index error messages, or

Explanation	other warnings that may need to be addressed to increase query response time.
	If there are no messages indicating any specific warnings, ask your administrator to increase the Core Query Timeout from 5 minutes to 10 minutes as described in the *System Security and User Management Guide*.

Message	`The session id is too large to be handled:<<eventId>`
Issue	The session id that you typed in, or got from the Legacy Events view or Navigate view is too large.
Explanation	If you manually typed the sessionId or edited a sessionId in the Events view, you may have created an integer that is too large for Events to process.

Behavior	When reconstructing network events with a large number of packets (>250) in the Events view > Packets panel, with the option to display only payloads enabled and the packets per page setting higher than the default (100), the current browser tab may become unresponsive for up to 45 seconds as it is working to render the payloads.
Issue	Depending on the amount of resources (memory and CPU) on the client machine and the number of packets in the event there may be a performance lag when displaying only payloads in packet reconstruction.
Explanation	To limit the amount of data processed in a reconstruction of a single event, change the Packets per Page setting in the footer to a lower value. 

Behavior	When working in the Version 11.4 Events view, the Query Profile drop-down menu and Column Group drop-down menu do not function.
Issue	You do not have permission to read columns groups and profiles. The default column group , Summary List, is applied to the Events list, and you cannot change the column group, create a column group, or delete a column group.
Explanation	This occurs only when the administrator has created a custom role for you instead of assigning the default Analyst role. Ask your administrator to enable column group read and profile read permission for your role.

Issue	No matching Endpoint data available on Investigate > Events view > Host tab.
Explanation	The Endpoint data may not be available due to any of the following: • **No Endpoint Deployment** – You must install Endpoint Log Hybrid, see "RSA NetWitness Endpoint" in the *Physical Host Installation Guide*.

- **Endpoint data is not captured for the host associated with the selected network event-** Make sure that NetWitness Endpoint Agent is installed, and expanded network visibility is configured to track the network events. To enable expanded network visibility, see “Creating Groups and Policies” in the *NetWitness Endpoint Configuration Guide*.

  Note: For Expanded Network Visibility to work, ensure the service user account used for aggregating Endpoint Log Decoder data to Endpoint Concentrator is assigned with the `decoder.manage` permission. For more information on how to assign roles and permissions, see "Services Security View - Aggregation Role" in the *Hosts and Services Getting Started Guide for RSA NetWitness Platform*.

- **Concentrators or Endpoint services are offline or very slow** – You must check the status (online or offline) of the services on Health and Wellness. If the service is online, you must check the Endpoint server logs and (Endpoint) Concentrator logs for details.
- **Endpoint data is rolled over for the host associated with the selected network event** - The Endpoint data may be rolled over due to data retention period configured. You must configure the data retention period to retain the endpoint data for a longer period. For more information, see “Configure Data Retention” in the *Data Privacy Management Guide*.

Investigate Reference Materials

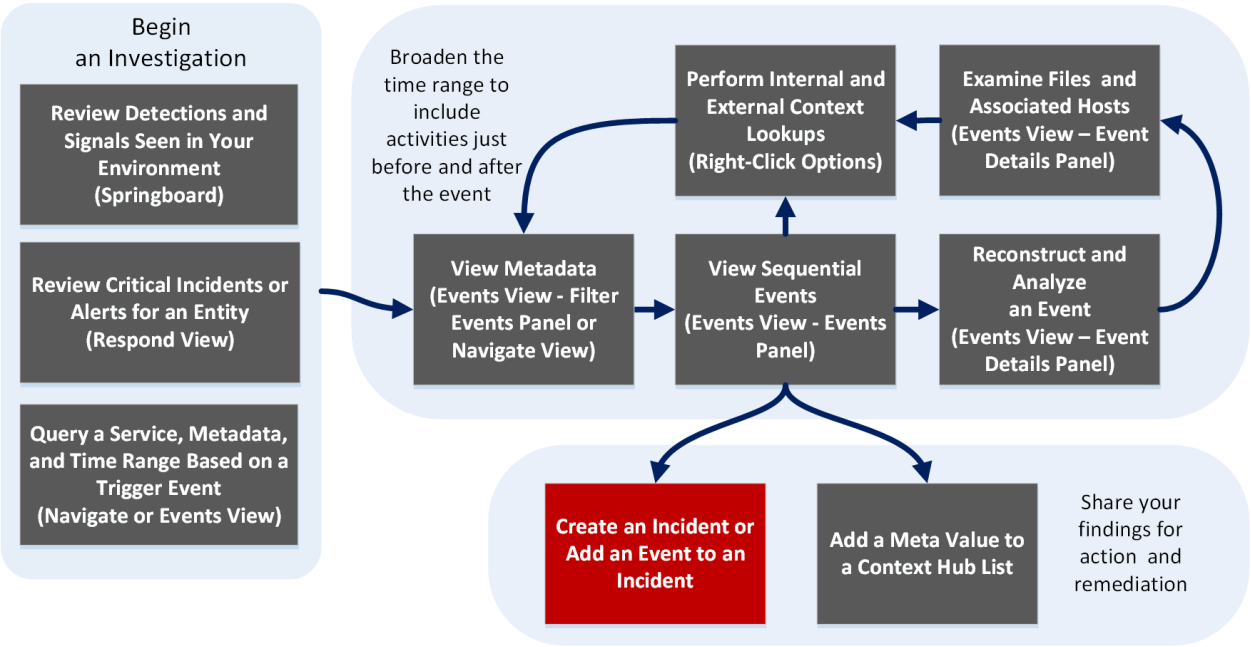
This section provides is intended to help you understand the purpose and application of NetWitness Investigate views. For each view, there is a brief introduction and a What Do You Want To Do table with links to related procedures. In addition some of the reference materials include workflows and Quick Looks to highlight important features in the user interface.

- Investigate View
- Navigate View
- Legacy Events View
- Events View
- Add/Remove from List Dialog
- Add Events to an Incident Dialog
- Column Groups Dialogs
- Context Lookup Panel
- Create an Incident Dialog
- Events View - Email Tab
- Events View - Text Tab
- Events View - Packet Tab
- Events View - File Tab
- Investigate Dialog
- Investigation Tab - User Preferences Panel
- Legacy Event Reconstruction View
- Manage Default Meta Keys Dialog
- Meta Groups Dialogs
- Navigate View
- Query Dialog
- Query Profiles Dialogs
- Settings Dialogs for Investigate Views

Add Events to an Incident Dialog

In the Add Events to an Incident dialog, analysts can add alerts to an existing incident so that incident responders look at the associated events as part of an incident response. To access this dialog while investigating a service in the Events view and the Legacy Events view, see Add Events to an Incident in the Events View and Add Events to an Incident in the Legacy Events View.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)

User Role	I want to ...	Show me how
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident*	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

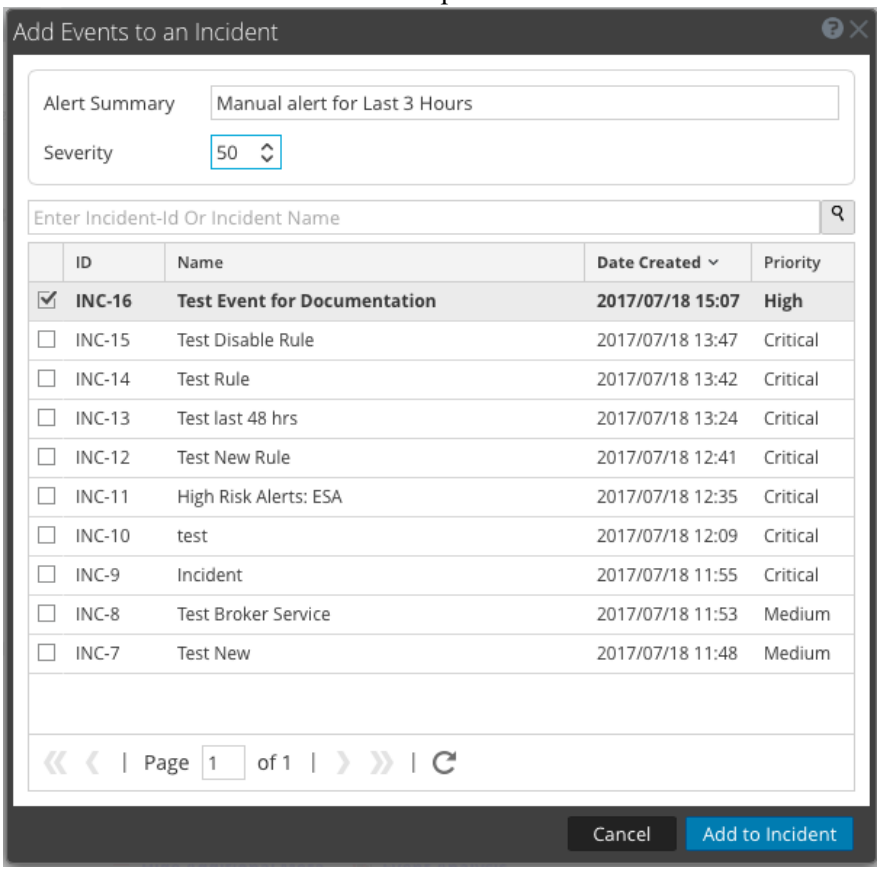
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View
- Legacy Events View

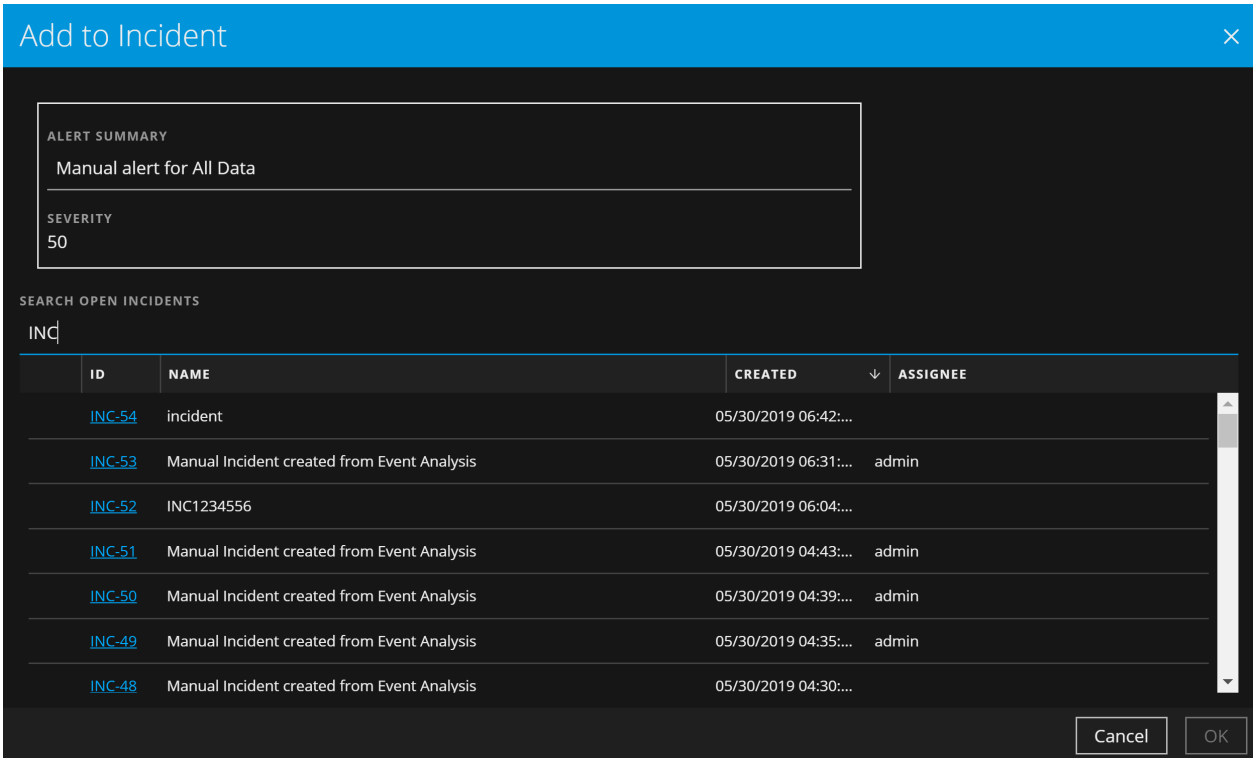
Quick Look

The following figure is an example of the Add Events to an Incident dialog in the Legacy Events. The table describes the information and options in the Add Events to an Incident dialog .



Feature	Description
Alert Summary	The Alert Summary field is filled by the query that produced the select alerts, which you selected to create this incident. The Severity field reflects the Severity of the selected alert, an integer between 1 and 100.
Search	Allows you to search for an existing event.
ID	The ID of the incident. You can sort IDs in ascending or descending order.
Name	The incident name. You can sort the Name in ascending or descending order.
Date Created	Displays the date and time the incident was created. You can sort the dates in ascending or descending order.
Priority	Displays the priority of the incident: either low or critical.
Cancel	Closes the dialog without saving changes.
Add to Incident	Adds the alerts to the incident. A dialog confirms that alerts are successfully added

The following figure is an example of the Add to Incident dialog in the Events view. The table describes the information and options in the Add to Incident dialog.



Feature	Description
Alert Summary	The Alert Summary field is filled by the query that produced the select alerts, which you selected to create this incident.
Severity	The Severity field reflects the Severity of the selected alert, an integer between 1 and 100.
Search Open Incidents	Allows you to search for an existing incidents.
ID	The ID of the incident.
Name	The incident name.
Created	Displays the date and time the incident was created.
Assignee	Displays the team member currently assigned to the incident
Cancel	Closes the dialog without saving changes.
OK	Adds the alerts to the incident. A confirmation message is displayed after the incident is successfully added

Add/Remove from List Dialog

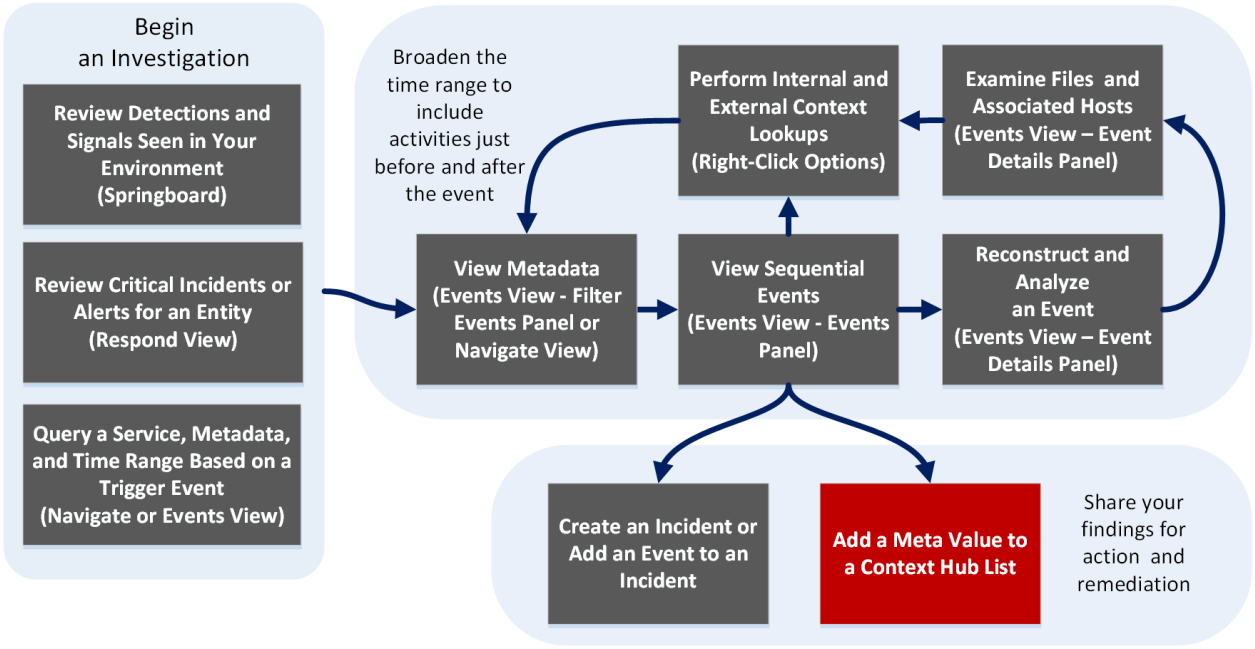
The Add/Remove from List dialog allows you to add an entity or meta value to an existing Context Hub list, remove an entity or meta value, or create a new Context Hub list containing the entity or meta value. When you look up an IP address or other entity and you find it suspicious or interesting, you can add it to a list that has been added as a data source. An example of a commonly used list is a white list or black list. This improves the visibility of the suspicious IP addresses and reduces false positives that do not need further investigation.

You can add entities or meta values to more than one list. For example, you can add them to one list for suspected domains related to command and control connections and to another list for Trojan connection IP addresses related to remote access. If a list is not available, you can create a list.

The dialog is available in NetWitness Investigate and in NetWitness Respond. When working in Investigate, in the Navigate view, Legacy Events view, or Events view, you can add meta values for the `Source IP`, `Destination IP`, or `Username` meta keys to an existing context hub list or you can create a new list containing the meta values. When you add meta values to a list, you can look up additional context on those meta values.

- To display the dialog in the Navigate view or the Legacy Events view, right-click a meta value under `Source IP`, `Destination IP`, or `Username`) and select **Add/Remove from List** in the context menu.
- To display the dialog in the Events view, hover over a value and select **Add/Remove from List** in the Actions section of the context tooltip.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

*You can perform this task in the current view.

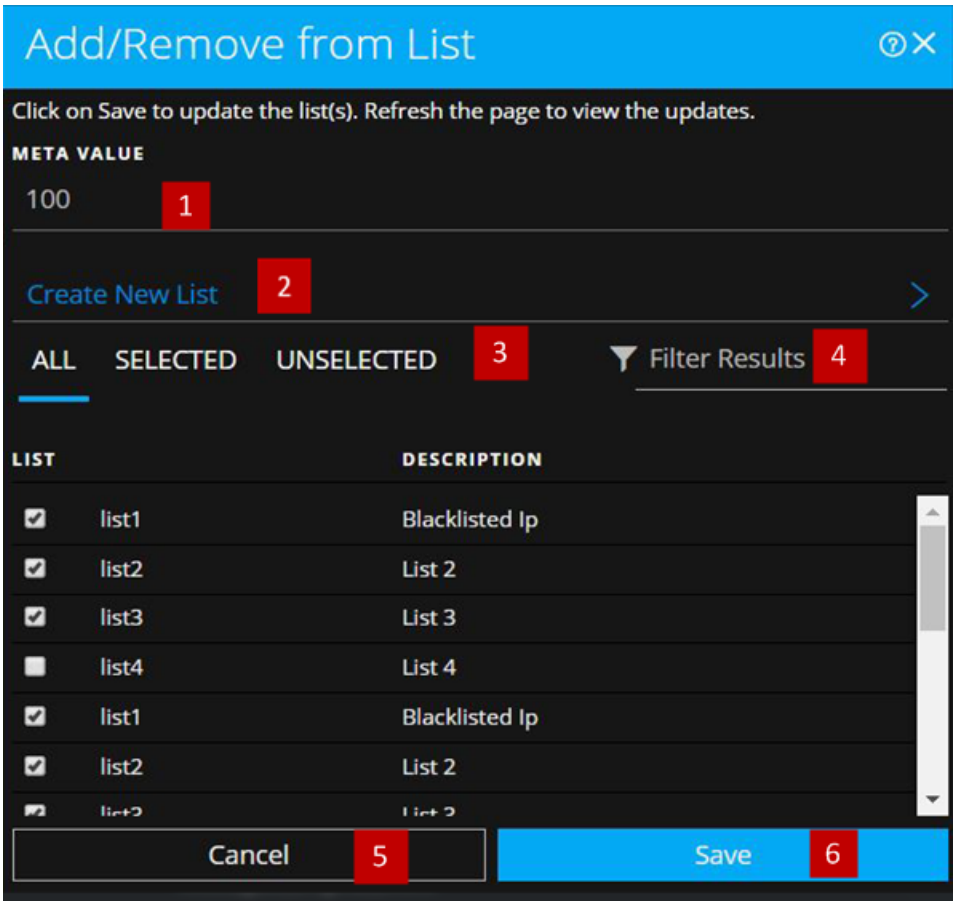
Related Topics

- Look Up Additional Context for Results
- Navigate View

- Legacy Events View
- Events View

Quick Look in the Events View

The following is an example of the **Add/Remove from List** dialog in the Events view.



1 Entities or meta values to be added or removed.
2 Create a new list using the selected meta.
3 Select any of the tabs: All, Selected, or Unselected.
4 Search using the list name or description.
5 Cancel the action.
6 Save to update lists or create a new list.

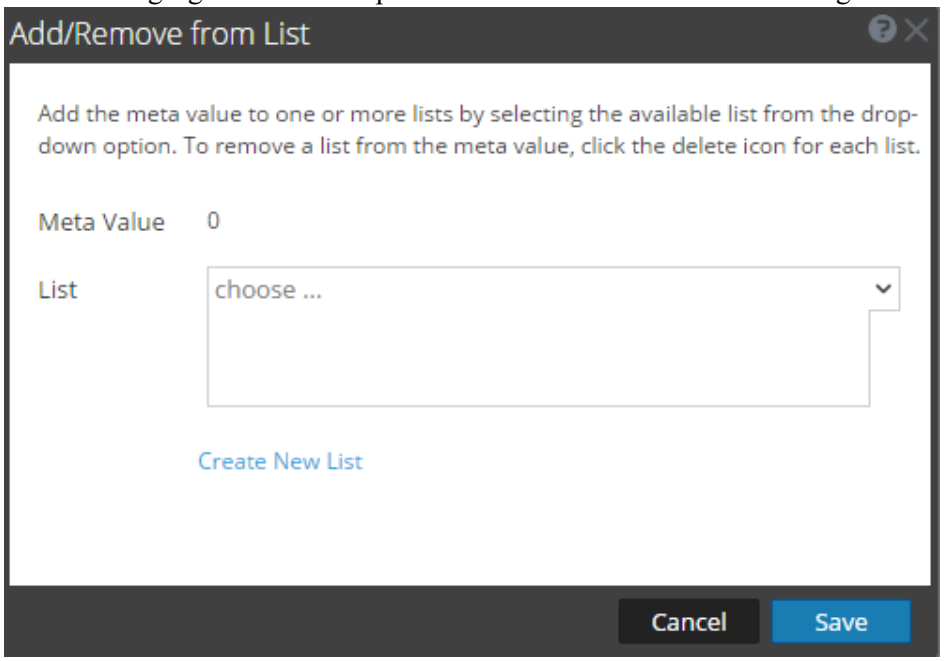
The following table shows the options in the Add/Remove from List dialog.

Option	Description
META VALUE	Displays the selected entity or meta value that needs to be added to or removed from one or more lists. You can also a create a new list using the selected value.

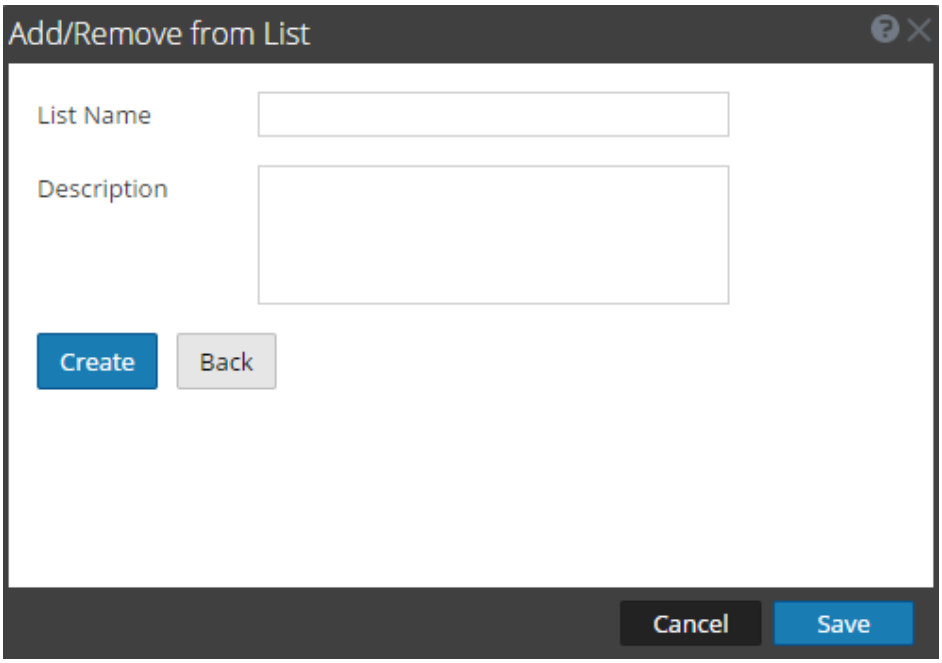
Option	Description
Create New List	Displays a dialog to create a new list using the selected meta value.
ALL	Shows all of the available Context Hub lists. The lists that contain the selected entity or meta value are selected. Select a checkbox to add an entity or meta value to a list. Clear a checkbox to remove it from the list.
SELECTED	Shows only the lists that contain the selected entity or meta value. (All lists are selected.)
UNSELECTED	Shows only the lists that do not contain the selected entity or meta value. (All lists are unselected.)
Filter Results	Enter the name or description of a specific list to search from multiple lists.
LIST	Displays the name of all the lists.
DESCRIPTION	Displays information about the selected list. The description that you provide when creating a list appears in this dialog. For example: This list contains all of the blacklisted IP addresses.
Cancel	Cancels the operation.
Save	Saves the changes.

Quick Look in the Navigate and Legacy Events Views

The following figure is an example of the Add/Remove from List dialog when initially opened.



The following figure shows the dialog when you select Create New List.



The following table describes the features of the Add/Remove from List and Create New List dialogs.

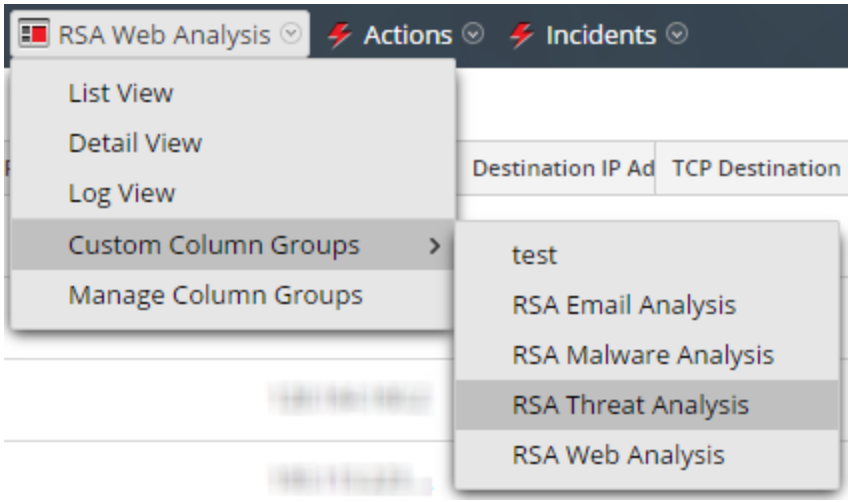
Feature	Description
Meta Value	The selected meta value to be added to the existing or new list.
List	The list to which the selected meta value must be added. A drop-down menu provides a list of available lists to which you can add the meta value.
Create New List	Opens a new dialog in which you can create a new list for the selected meta value.
List Name	The name of the new list.
Description	The description of the new list.
Create	Creates a new list after entering the required fields.
Back	In the new list mode, cancels the new list creation and returns to the original dialog.
Cancel	Cancels the addition of the meta value to a list and closes the dialog.
Save	Saves the changes made to the lists and closes the dialog.

Column Groups Dialogs

Column groups allow you to format the events list to include only the relevant meta keys in the Events view and Legacy Events view (see Use Columns and Column Groups in the Events List). When the events list in Investigate is populated with events, each column lists the values returned for a meta key. Changing the meta keys displayed in the events list is a useful method of narrowing the focus of your investigation. For example, the default column group includes columns for **Collection Time**, **Type**, **Theme**, **Size**, and **Summary**. These are just the basic information, not specialized in any way. The RSA Email Analysis list has only that contain information useful when investigating email.

The column group definition includes the meta keys to use as column titles, the position of the column in the list, and the default width of the column. You can add, delete, import, export, and edit column groups. At fresh installation, built-in column groups are available. The built-in column groups are prefixed with RSA and can be duplicated but cannot be edited or deleted. You can also create custom column groups.

- The Create Column Group dialog is for the 11.4 and later Events view. To access this dialog, select **Column Group > New Column Group** in the **Events** view toolbar.
- The Column Group Details dialog is for the 11.4 and later Events view. To access this dialog, select **Column Group** in the **Events** view toolbar, then click the edit icon ( ) next to a custom column group name.
- The Manage Column Groups dialog is for the Legacy Events view (Version 11.4) , and the Events view (Version 11.4 and earlier). The Manage Column Groups dialog has these features that are not yet available in the Create Column Group dialog: set column width, import, and export. To access this dialog, go to **Investigate > Legacy Events** and in the **View** drop-down list select **Manage Column Groups**. The **View** option is named for the current value, for example, Detail View, List View, Log View, or the currently selected column group.



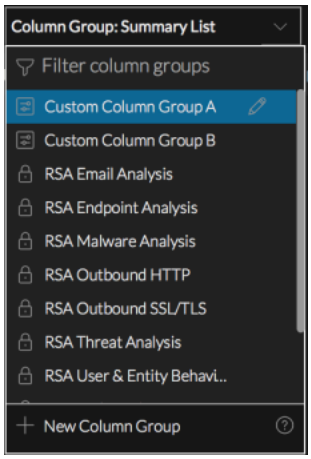
After column groups are defined, you can use them in other Investigate views. In the Navigate view, Query Profiles allow you to select a column group to use when the profile is applied. In the Events view and the Legacy Events view, you can select a column group to apply to the Events panel.

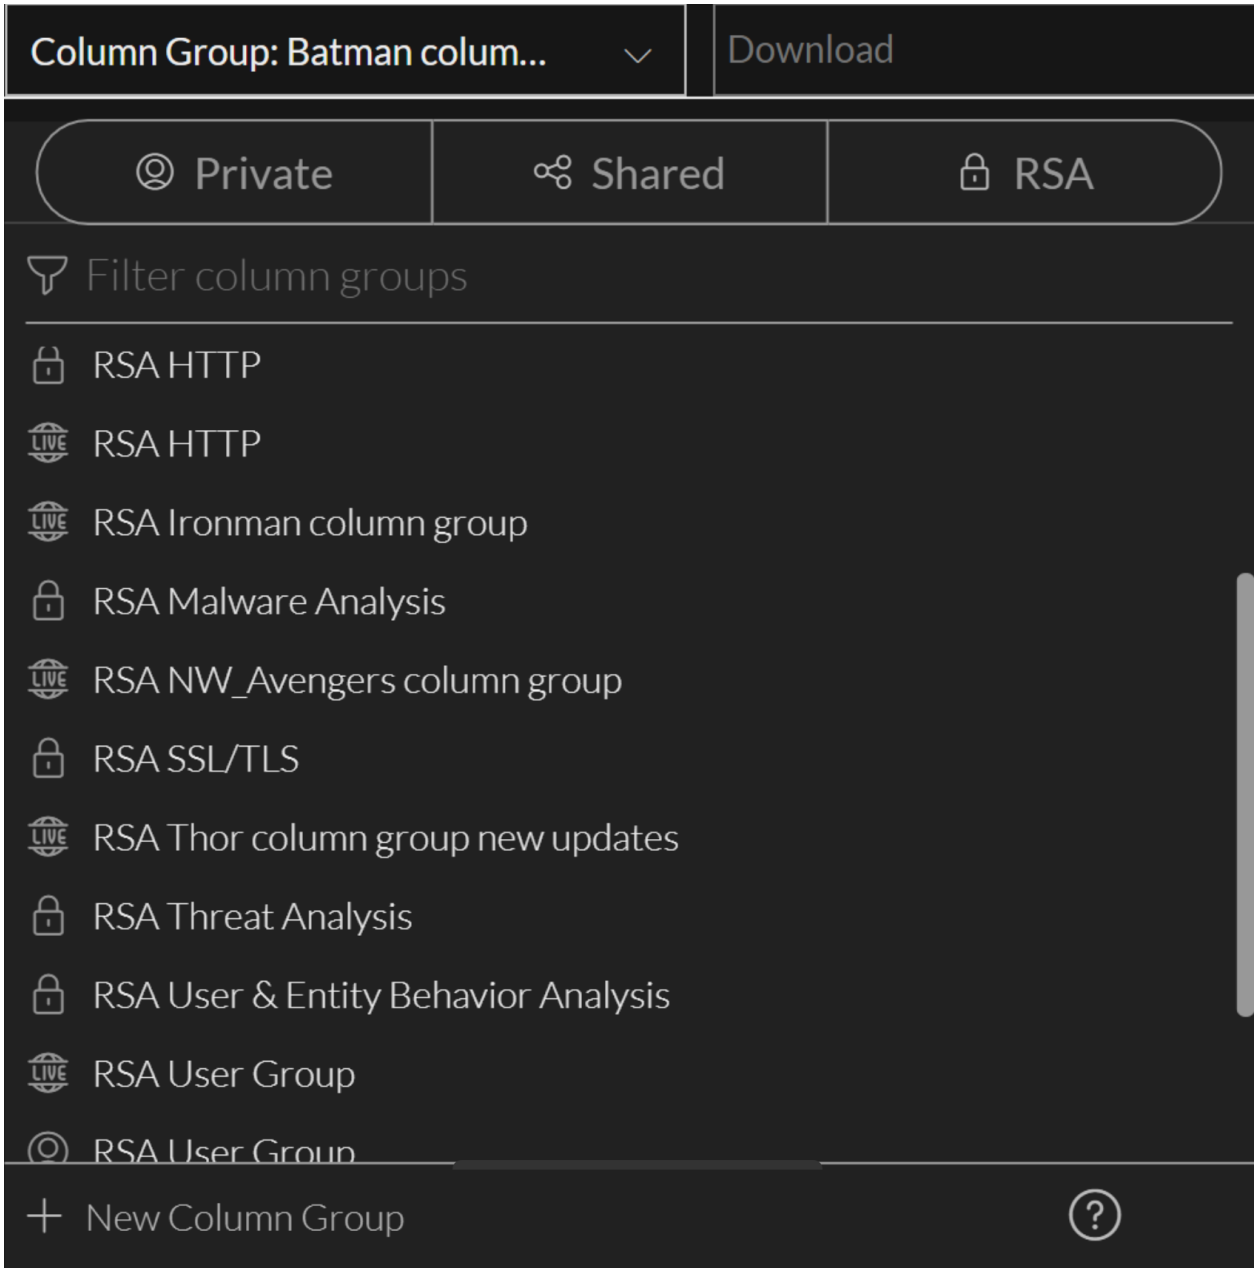
Related Topics

- How NetWitness Investigate Works
- Events View
- Legacy Events View

Quick Look - Column Group Menu, Create Column Group Dialog, and Column Group Details Dialog

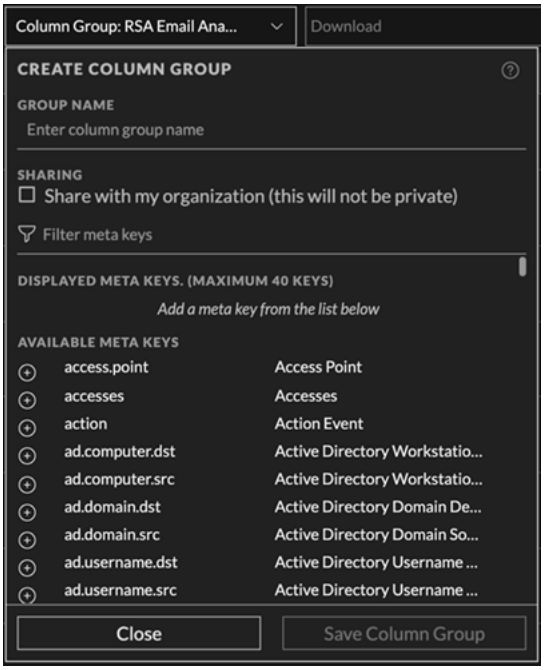
This section introduces the Column Group Menu, Create Column dialog, and the Column Group Details dialog. The following figures are examples of the Column Group menu. The example on the left (Version 11.4) has a custom column group highlighted so that the edit icon is visible. The example on the right (Version 11.5) has a built-in column group highlighted so that the information icon is visible. The table describes the options.

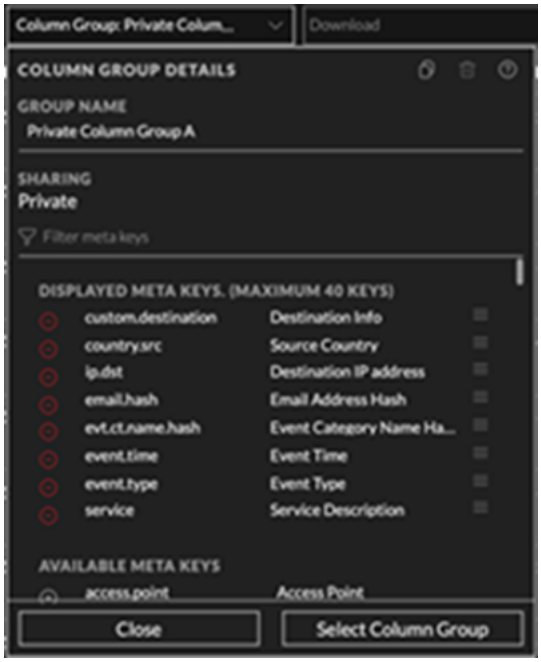

Column Group: Summary List
Filter column groups
Custom Column Group A
Custom Column Group B
RSA Email Analysis
RSA Endpoint Analysis
RSA Malware Analysis
RSA Outbound HTTP
RSA Outbound SSL/TLS
RSA Threat Analysis
RSA User & Entity Behavi...
New Column Group


Column Group: Batman colum...
Download
Private
Shared
RSA
Filter column groups
RSA HTTP
RSA HTTP
RSA Ironman column group
RSA Malware Analysis
RSA NW_Avengers column group
RSA SSL/TLS
RSA Thor column group new updates
RSA Threat Analysis
RSA User & Entity Behavior Analysis
RSA User Group
RSA User Group
New Column Group

Feature	Description
Visibility Options	(Version 11.5 and later) Control the types of column groups that are visible in the list, using any combination of the visibility options (blue = selected, black = not selected): **Private** = display private groups that only you can manage **Shared** = display shared groups that anyone in your organization can manage **RSA** = display built-in groups that only RSA can manage The visibility options work together with the Filter Column Groups field. If the visibility option is hiding built-in groups (which include "RSA" in the group name) and you search for a name that contains "RSA," the list is empty.
Filter Column Groups	Filters the list of column groups as you type text so that only group names that contain that text are displayed.
Column Group List	The list of column groups consists of custom and built-in groups, which are distinguished by the icons that precede the name. In Version 11.5 and later, custom groups can be shared or private. The RSA column groups are built-in column groups. The icons distinguish the private custom groups, shared custom groups, and built-in groups.
New Column Group	Displays the Create Column Group dialog, where you can create a custom column group.

The Create Column Group dialog, shown in the figure on the left, allows you to define a custom column group. The figure on the right illustrates the Column Group Details dialog, in which you can edit a custom column group. The table describes the fields and options in the dialogs.

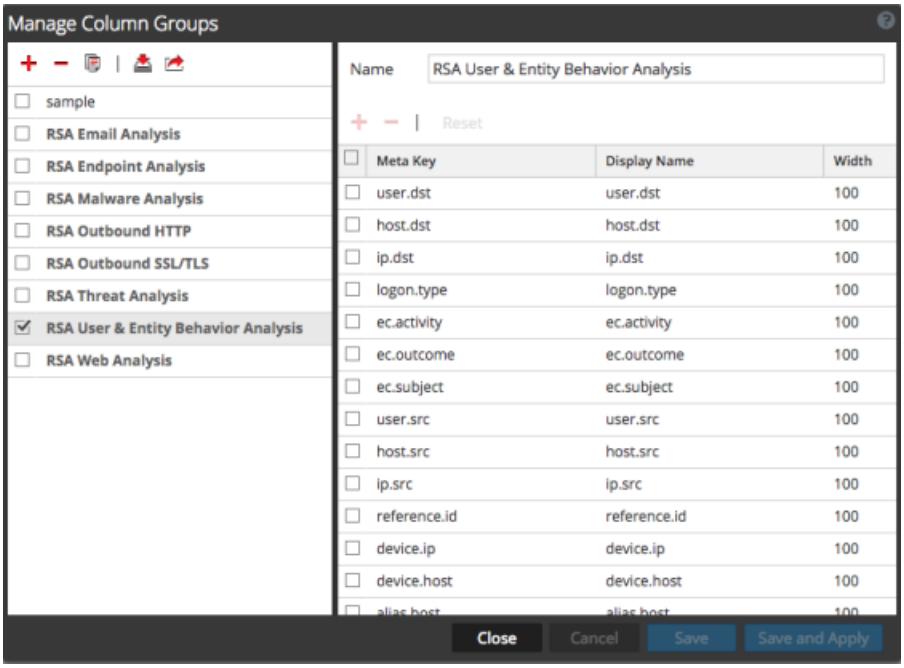




Feature	Description

Feature	Description
	Deletes the custom column group in the Column Group Details dialog. This action is irreversible and applies globally; the column group is no longer available to anyone who is using it on this service.
Group Name	Displays the name of the column group. The name must be unique and contain fewer than 64 characters. You can type in this field to edit the name in a custom column group.
Sharing	In Version 11.5 and later, you can create column groups that are shared or private. This setting is available when you first create the group. After it is created, you cannot change a shared column group to private, or a private column group to shared.
Filter Meta Keys	Filters the Displayed Meta Keys and Available Meta Keys based on the text that you type. Only meta keys that contain the typed text are displayed.
Displayed Meta Keys	Displays a scrollable list of meta keys that are selected for use in the custom column group. You can add meta keys in the Available Meta Keys list to this list, remove meta keys from this list  , and drag meta keys up or down to change the order in this list ().
Available Meta Keys	Displays a scrollable list of meta keys that are available (on the service) for use in the custom column group. You can add them to the Available Meta Keys list. Clicking  next to the meta key name adds it to the Displayed Meta Keys list.
Close button	Closes the dialog.
Save Column Group	For the Create Column Group dialog only, saves the new column group.
Reset	For the Column Group Details dialog only, reverts the edited column group to the last saved state.
Update Column Group	For the Column Group Details dialog only, applies changes to an edited column group.
Select Column Group	Applies the column group.

Quick Look - Manage Column Groups Dialog



The Manage Column Groups dialog has two panels: Groups and Settings. At the bottom of this dialog are four buttons: Close, Cancel, Save, and Save and Apply.

The left panel is the Groups panel. This is where you can add, delete, import, or export column groups. At the top of the panel is a toolbar. Below the toolbar is a list of added column groups, where you can select one or more groups.

The following table lists the actions in the toolbar.

Action	Description
+	Adds a column group. Clicking this button highlights the Settings panel on the right, where you can name the column group and add or delete meta keys. At least one meta key is required to add a group.
−	Deletes a column group. A confirmation dialog is displayed before the selected group is deleted. OOTB column groups cannot be deleted.
	Creates a copy of the selected column group.
	Displays the Import Column Groups dialog, where you can select a file to upload.
	Exports one or more selected groups to your local file system.

The right panel is the Settings panel. This is where you can create and edit column groups. This panel contains the Name field, a toolbar, and a list. The following table describes the features of the Settings panel.

Feature	Description
Name	The name of the selected column group.
	Adds a new row to the list of meta keys, where you can open a drop-down menu to select a new meta key.
	Deletes one or more selected meta keys. Displays a confirmation dialog before deleting.
Reset	Returns the column group to its most recently saved settings.
Meta Key	Lists the meta keys added to the selected column group.
Display Name	Lists the names of the meta keys as they are displayed in the Navigate, Events, and Event Analysis views.
Width	Specifies the width of each meta key's column. The width can be set between **10** and **1000**. The default width is **100**.

The following table provides descriptions of the action buttons.

Feature	Description
Close	Closes the dialog without saving.
Cancel	Cancels all unsaved changes.
Save	Saves all changes without closing the dialog.
Save and Apply	Saves and applies all changes immediately, closing the dialog.

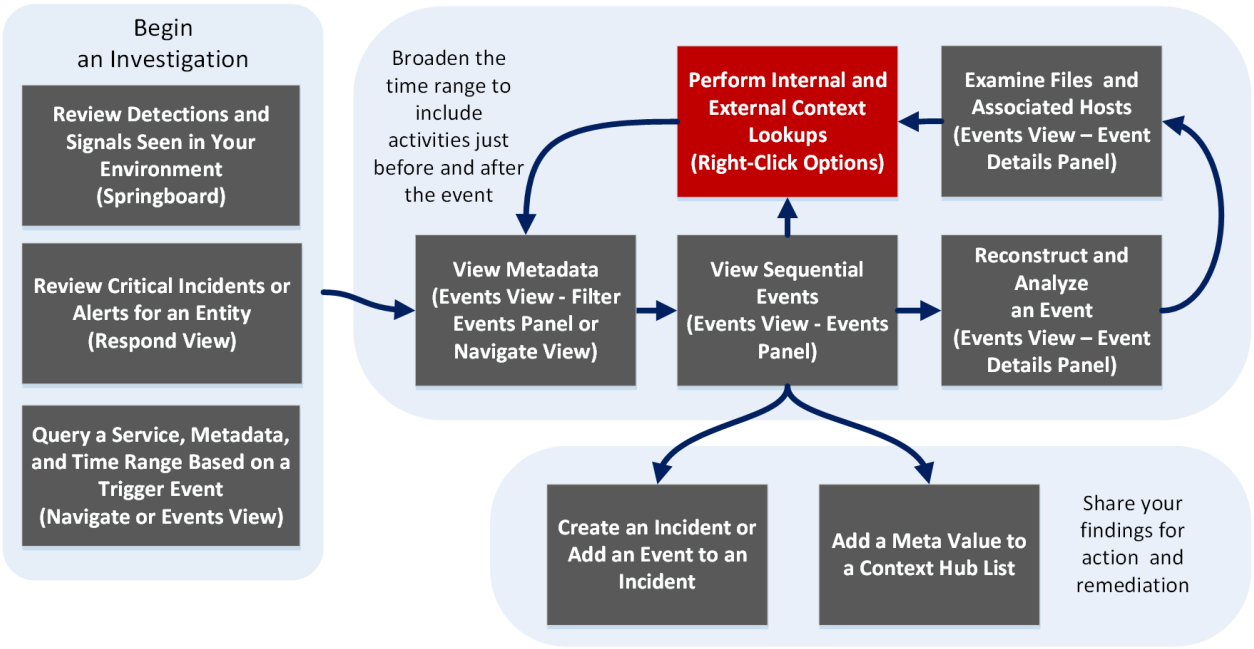
Context Lookup Panel

After an administrator configures the Context Hub service, you can view the contextual information for the meta values in the Navigate view, Legacy Events view, and Events view (Version 11.2 and later). The Context Hub service is pre-configured with a default meta type and meta key mapping. For information about mapping of the context hub meta value with investigation meta key, see "Manage Meta Type and Meta Key Mapping" in the *Context Hub Configuration Guide*.

The Context Lookup panel is displayed on the right side of the Navigate view, Legacy Events view, or Events view. Meta values that have been added to a Context Hub list are highlighted in gray in the Navigate view or Legacy Events view results. In the Events view, they are marked by an underscore. When you right-click a highlighted value and select **Context Lookup** in the resulting context menu, the lookup results are displayed in the Context Lookup panel for configured sources for the selected meta value. You can select a source in the Context Lookup panel icon bar to view the contextual information.

There are some differences between the appearance and contents of the Context Lookup panel when open in the Navigate view or Events view and when open in the Events view.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*

User Role	I want to ...	Show me how
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

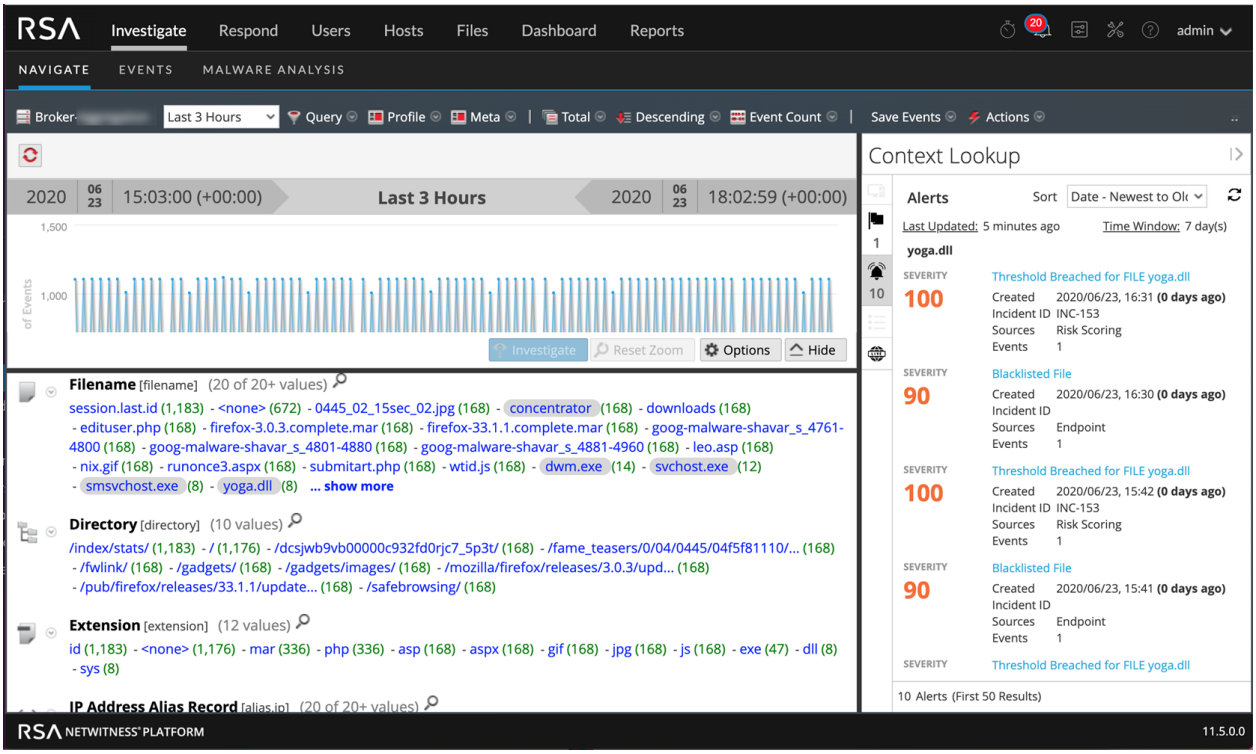
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Legacy Events View
- Navigate View
- Events View
- "NetWitness Feedback and Data Sharing" in the *Live Services Management Guide*

Quick Look (in the Navigate and Legacy Events Views)

The following figure is an example of the Context Lookup panel as it appears in the Navigate view. Controls and features are described in the table.



Feature	Description
Source Options Bar	Displays the icons for the available sources: Endpoint, Incidents, Alerts, and Lists.
Source Name	Displays the source name based on the selected icon: • Endpoint • Incidents • Alerts • Lists • Live Connect
Sort	Provides a drop-down of sort options for the listed context information. Possible sort options are Severity - High to Low, Severity Low to High, Date - Oldest to Newest. and Date - Newest to Oldest. The sorting options vary by source type.
	Refreshes the lookup results.

Feature	Description
<n items> (First <n> Results)	The footer provides a count of results currently displayed and the total number of results. For example, 5 Alerts (First 50 Results).

Incidents

Incidents are displayed based on time first (Newest to Oldest) and then priority status. The following information is displayed for incident lookups:

- Incident Name and ID
- Priority status of the incidents
- Risk Score value of the incidents
- Date when the incident was created
- Status of the incident
- Assignee for the incident
- Last Updated: Indicates when contextual data was last fetched from data source and updated to cache.
- Time window: This is based on the value that is set for the "Query Last (Days)" field in the Configure Respond window. For details, see the "Configure Respond as a Data Source" topic in the *Context Hub Configuration Guide*.
- Sort: This drop-down field provides options to change the sorting of result based on time or priority.

Alerts

Alerts are displayed based on the Severity. ;The following information is displayed for alert lookups:

- Alert Name
- Severity value of the alerts
- Date when the alert was created
- Incident ID: This is the ID of the incident that the alert is associated with (If any).
- Sources: Event source name
- Number of events associated with the alert.
- Last Updated: Indicates when contextual data was last fetched from data source and updated to cache.
- Time window: This is based on the value that is set for the "Query Last (Days)" field in the Configure Respond window. For details, see the "Configure Respond as a Data Source" topic in the *Context Hub Configuration Guide*
- Sort: This drop-down field provides option to change the sorting of result based on time or priority.

Lists

The following information is displayed for list lookups.

- List Name
- Owner who created the list
- Created Date
- Last Updated Date
- Description of the list

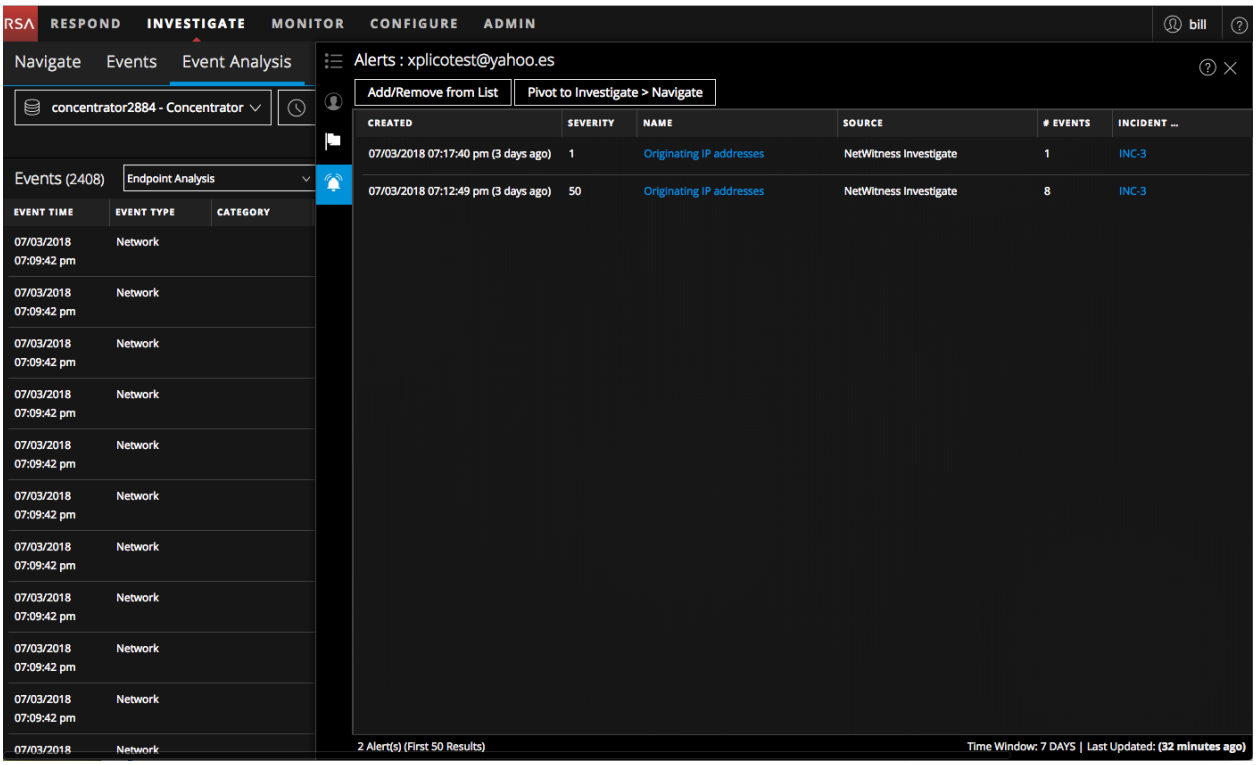
Endpoint

The following information is displayed for Endpoint lookups.

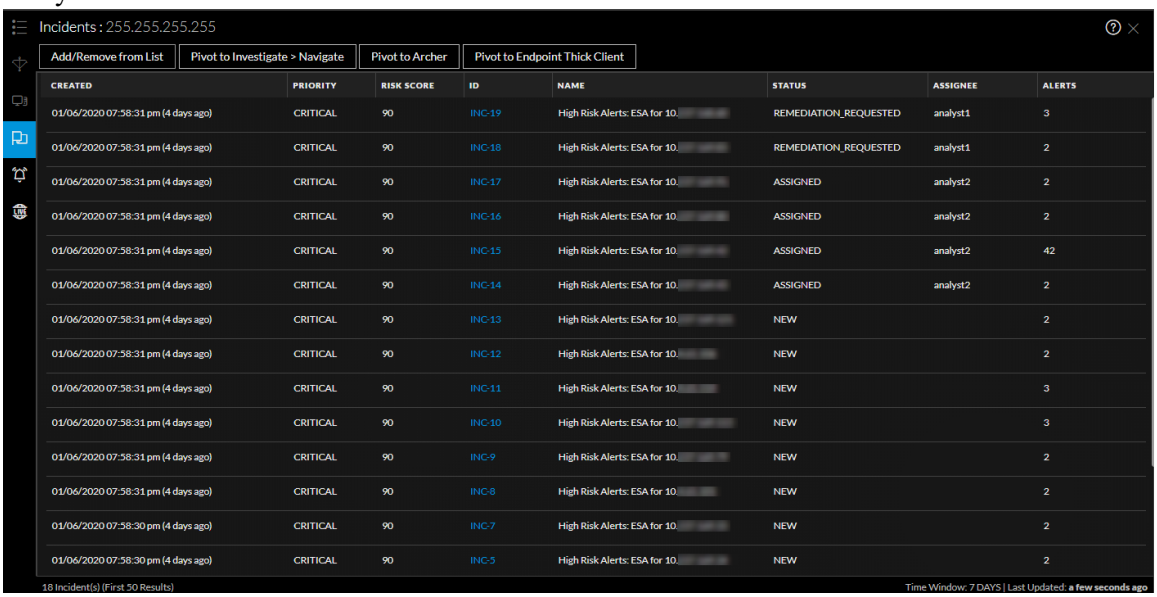
- Machine name and IP address of the machine.
  By clicking on the IP or Endpoint machine name, you will be navigated to Endpoint UI to perform further investigation.
- Last Updated: Indicates when contextual data was last fetched from data source and updated to cache.
- Machine Score: A machine IIOC score is aggregated based on the module scores.
- Number of modules: Number of active files for the selected machine.
- Last Updated: Indicates when the scan results were last updated in Endpoint database.
- Last Login User
- Machine MAC Address
- Operating System Version
- Admin Notes (if any)
- Admin Status (if any)
- Top Suspicious Modules (Modules that have an IIOC score > 500). This is based on the value set for "Minimum IIOC Score" field in the Configure Endpoint window. The default value for "Minimum IIOC Score" is 500.
- Machine IIOC Levels

Quick Look in the Events View (Version 11.2 and Later)

The following figure is an example of the Context Lookup panel as it appears in the Events view.



The contextual information or query results displayed in the Context Lookup panel depends on the selected entity and the associated data sources. The Context Lookup panel has separate tabs for each of the data sources. The tabs are: List data source, Archer, Active Directory, Endpoint, Incidents, Alerts, Live Connect, and REST API. The following figure shows the Context Lookup panel for a selected entity in the Incident Details view.



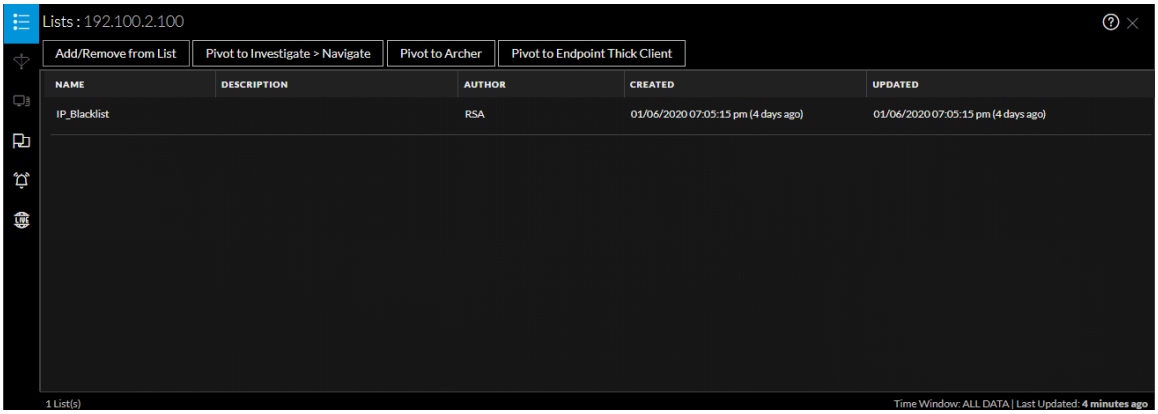
The following table describes the data available on each tab and the supported entities.

Tab	Description	Supported Entities
 (Lists)	Displays all of the list data associated with the selected entity or meta value. The result is sorted by the last updated list.	All entities
 (Archer)	Displays asset information along with criticality ratings using the Archer data source.	IP, Host, and Mac
 (Active Directory)	Displays all user information for the selected user.	User
 (NetWitness Endpoint)	Displays the NetWitness Endpoint data source information for the selected entity or meta value, which includes the Machines, Modules, and IIOC levels. Modules are by highest IOC score to lowest IIOC score and IIOC levels are sorted by highest IOC levels to lowest IOC levels.	IP, MAC address, and Host
 (Incidents)	Displays the list of incidents associated with the selected entity or meta value. The result is sorted by newest incidents to oldest incidents.	All entities
 (Alerts)	Displays the list of alerts associated with the selected entity or meta value. The result is sorted by newest alerts to oldest alerts.	All entities
 (Live Connect)	Displays information related to Live Connect.	IP, Domain, and Filehash
 (File Reputation)	Displays file reputation status for Filehash entities.	Filehash entities

Tab	Description	Supported Entities
TI	Displays information for STIX data sources.	IP address, email address, domain, filename, URL's, and file hash. **Note:** The context lookup for email address and URL will be displayed only if these metas are mapped. Navigate to (Admin) > System > Investigation > Context Lookup.
REST API	Displays the list of REST APIs (enabled in Context Hub) associated with selected the entity.	All entities

Lists Tab

The Context Lookup panel for Lists shows one or more lists associated with the selected entity or meta value. The following figure is an example of the Context Panel for Lists, and the table describes the fields.



Field	Description
Name	The name of the list (defined while creating the list).
Description	The description of the list (defined while creating the list).
Author	The owner who created the list.
Created	The date when the list was created.
Updated	The date when the list was last updated or modifed.
Count	The number of lists in which the selected entity or meta value is available.
Time Window	The time window based on the value set for the "Query Last" field in the Configure Responses dialog. By default, all Lists data is fetched.
Last Updated	The time when Context Hub fetched and stored the lookup data in cache.

Archer Tab

The Context Lookup panel for Archer displays asset information along with criticality ratings using the Archer data source for IP, Host, and Mac entities. The following figure is an example of the Context Lookup panel for Archer, and the table describes each field.



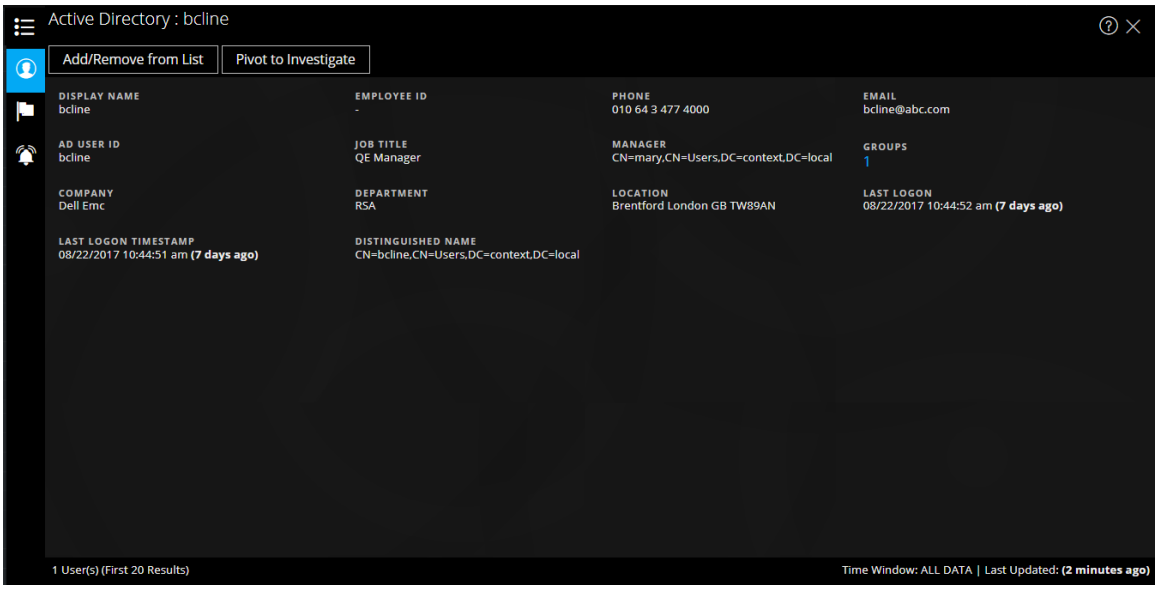
Field	Description
Criticality Rating	The device operational criticality based on the applications it supports. The criticality ratings can be set as Not Rated, Low, Medium-Low, Medium, Medium-High, or High.
Risk Rating	The calculated risk rating for the device based on the most recent assessment and the average risk rating of facilities using the device. The risk rating can be set as Severe, High, Medium, Low, or Minimal.
Device Name	The unique name of the device.
Host Name	The host name of the device.
IP Address	The primary internal IP address of the device.
Device ID	The automatically populated value that uniquely identifies the record across all applications within the system.
Type	The device type, for example, server, laptop, desktop, and others.
Facilities	Links to records in the Facilities application that are related to this device.
Business Unit	Links to records in the Business Unit application that are related to this device. For more than three business unit values, you can hover over the field to view the values.
Device Owner	The person who is responsible for the device and receives read and update rights of the record.

Field	Description
Count	The number of assets available.
Time Window	The time window based on the value that is set for the "Query Last" field in the Configure Responses dialog. By default, all data for Archer is fetched.
Last Updated	The time when Context Hub fetched and stored the lookup data in cache.

Note: In the localized versions, only these twelve fields are displayed: Criticality Rating, Risk Rating, Device Owner, Business Unit, Host Name, MAC Address, Facilities, IP Address, Type, Device ID, Device Name, and Business Processes.

Active Directory Tab

The following figure is an example of a Context Lookup panel for Active Directory.



The Context Lookup panel for Active Directory displays all the related information, incidents, and alerts for a user. You can perform a look up using the following formats:

- userPrincipalName
- Domain\UserName
- sAMAccountName

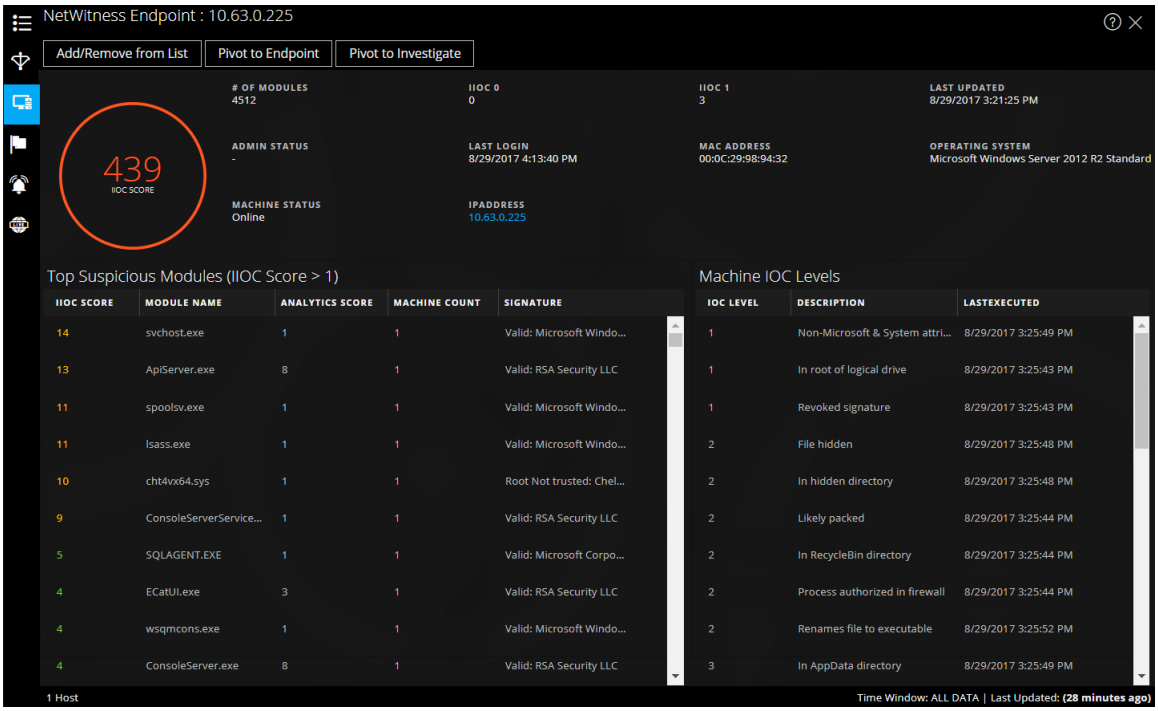
The following information is displayed for Active Directory.

Field	Description
Display Name	The name of the user.
Employee ID	The employee ID of the user.

Field	Description
Phone	The phone number of the user.
Email	The email ID of the user.
AD User ID	The unique identification of the user within an organization.
Job Title	The designation of the user.
Manager	The name of the user's manager.
Groups	The list of groups of which the user is a member.
Company	The name of the user's company.
Department	The department name to which the user belongs within the organization.
Location	The location of the user.
Last Logon	The time when the user logged into the system, only if the Global Catalogue is defined.
Last Logon TimeStamp	The time when the user logged into the system.
Distinguished Name	The unique name assigned to the user.
Count	The number of users.
Time Window	The time window based on the value that is set for the "Query Last" field in the Configure Data Source Settings dialog. By default, all data for Active Directory is fetched.
Last Updated	The time when Context Hub fetched and stored the lookup data in cache.

NetWitness Endpoint Tab

The following figure is an example of the Context Lookup panel for NetWitness Endpoint.



The following information displayed for IIOCs.

Field	Description
# Of Modules	The number modules that are looked up.
Admin Status	The admin status (if any).
Last Updated	The time when the data was last refreshed.
Last Login	The time when the user last logged in.
MAC Address	The Machine MAC Address.
Operating System	The Version of the Operating System used by the NetWitness Endpoint machine.
Machine Status	The state of the module being viewed: Online, Offline, Active, or Inactive.
IP Address	The IP address of the specific module.

The following information is displayed for modules.

Field	Description
IIOC Score	A machine IIOC score is an aggregated score based on the module scores. This is based on the value set for Minimum IIOC Score field in the Context Hub Data Source Settings dialog. The default value for Minimum IIOC Score is 500. See "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.

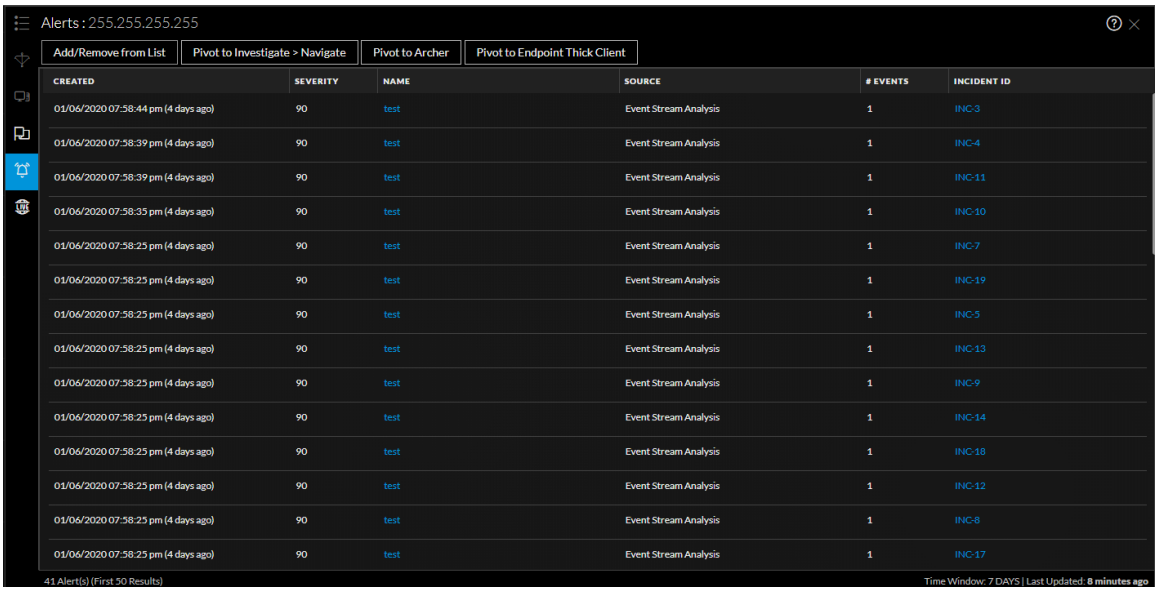
Field	Description
Module Name	The name of the module that is being looked up.
Analytic Score	The number of active files for the selected machine.
Machine Count	The number of machines on which that particular IOC got triggered.
Signature	Indicator of whether the file is signed or unsigned, valid or invalid, and signatory information. For example, Google, Apple, and so on.

The following information is displayed for machines.

Field	Description
IOC Levels	The IOC levels.
Description	The description for he IOC level if available.
Last executed	The time when the action was executed.
Count	The number of hosts that are being looked up.
Time Window	The time window based on the value set for the Query Last field in the Configure Data Source Settings dialog. By default, all data for NetWitness Endpoint is fetched.
Last Updated	The time when scan results were last updated in NetWitness Endpoint database.

Alerts Tab

The following figure is an example of Context Panel for Alerts that is displayed based on time first (Newest to Oldest) and then severity.

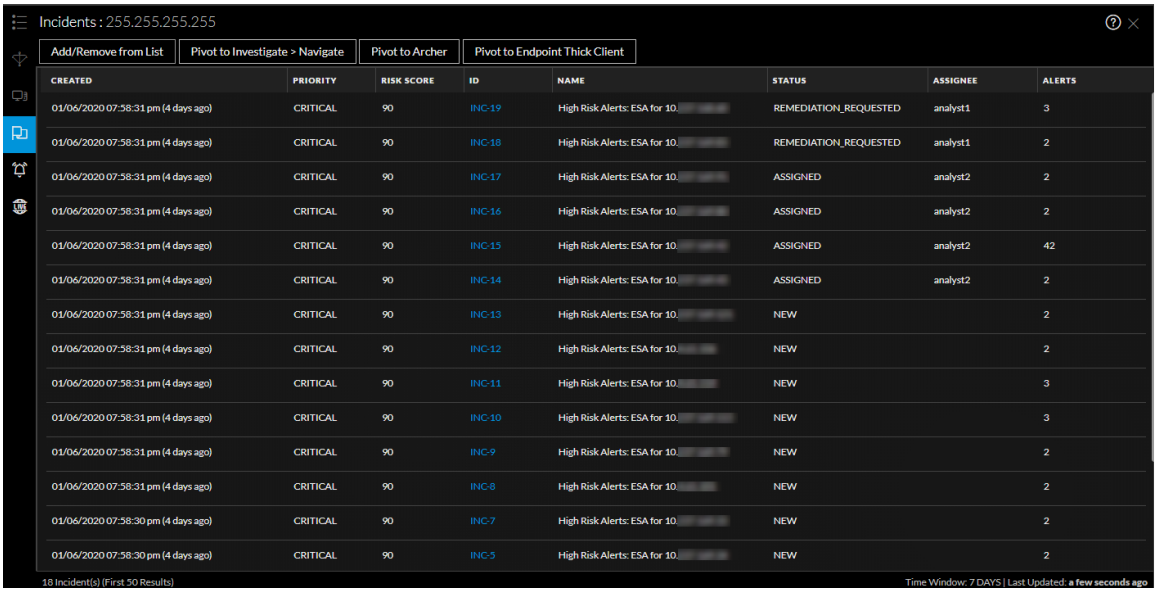


The following information is displayed in the Context Lookup panel for Alerts.

Field	Description
Created	The date and time when the alert was created.
Severity	The severity value of the alerts.
Name	The name of the alert. You can click the name to view the details of a specific alert.
Source	The alert source name from which the alert is triggered.
#Events	The number of events associated with the alert.
Incident ID	The ID of the incident (if any) with which the alert is associated. You can click the ID to view the details of a specific alert.
Count	The number of alerts. By default only the first 100 alerts are displayed. For more information on how to configure the settings, see "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.
Time Window	The time window based on the value set for the Query Last field in the Configure Data Source Settings dialog. By default, the alert data for last 7 days is fetched.
Last Updated	The time when contextual data was last fetched from data source.

Incidents Tab

The following figure is an example of the Context Panel for Incidents, which is based on time first (Newest to Oldest) and then priority status.

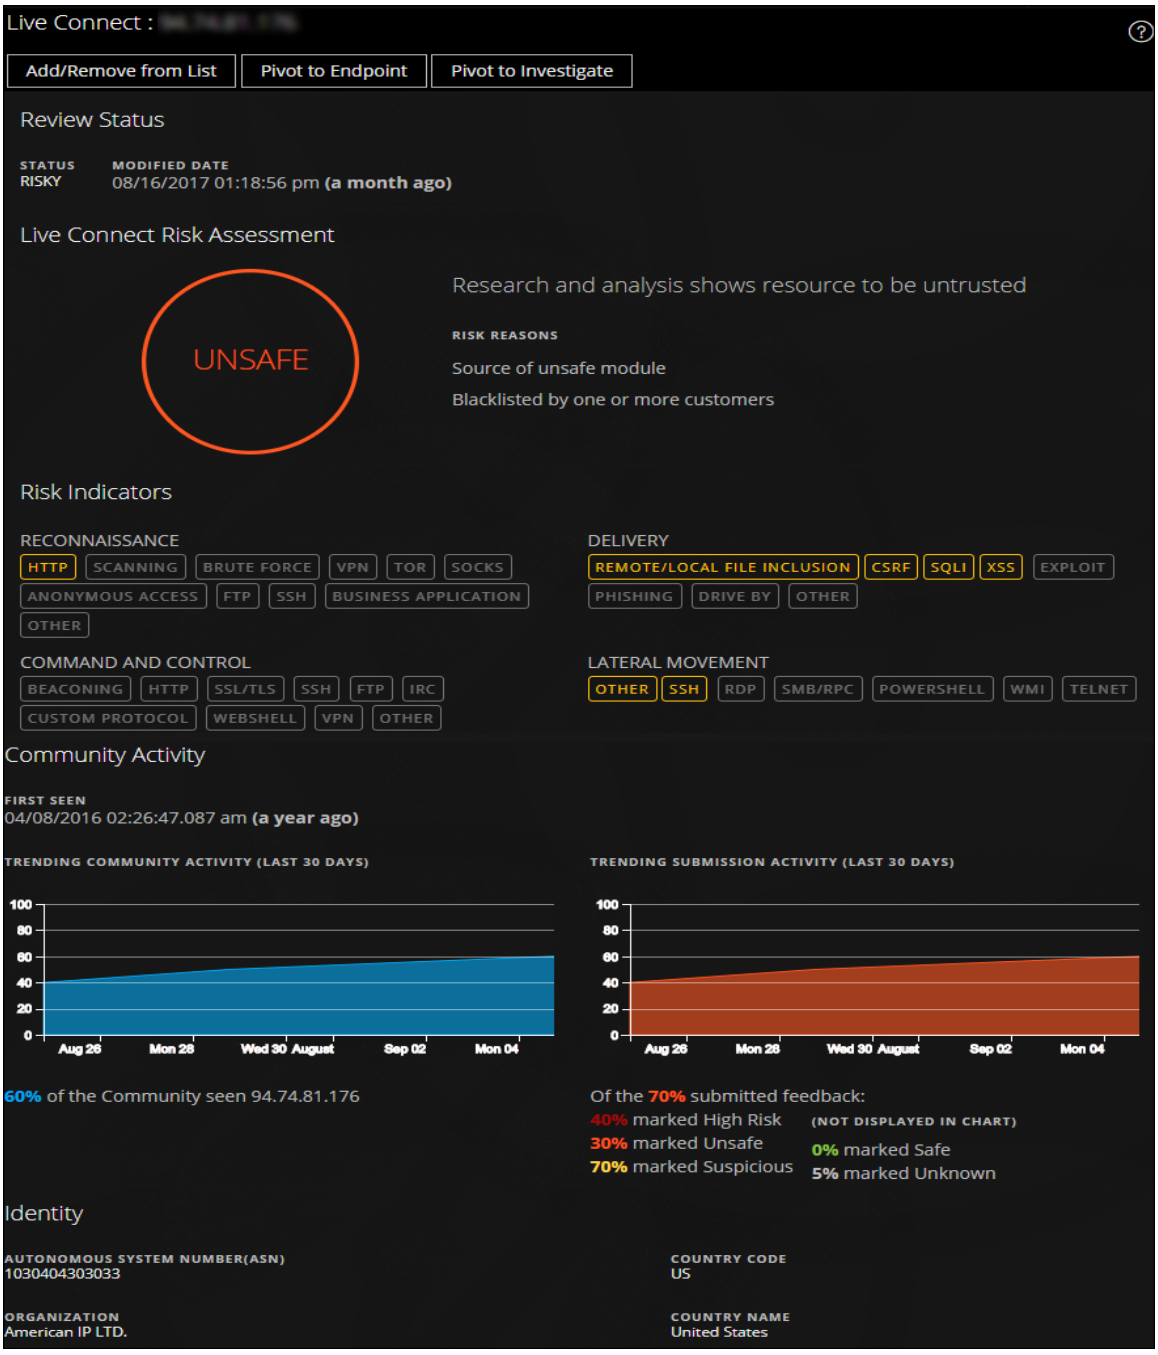


The following information is displayed in the Context Lookup panel for Incidents.

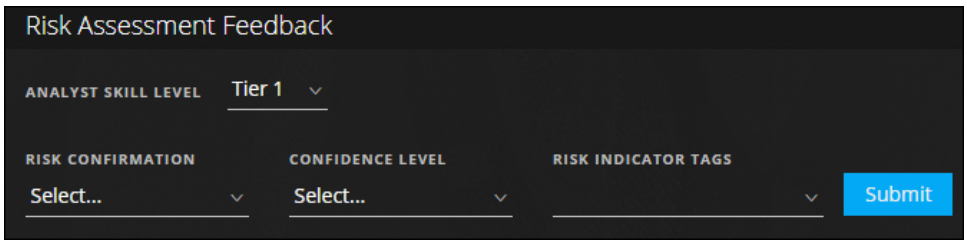
Field	Description
Created	The date when the incident was created.
Priority	The priority status of the incidents.
Risk Score	The risk score of the incidents.
ID	The Incident ID of the incident. You can click on the ID to display further details about the incident.
Name	The incident name.
Status	The status of the incident
Assignee	The current owner of the incident.
Alerts	The number of alerts associated with the incident.
Count	The number of incidents. By default only the first 100 incidents are displayed. For more information on how configure the settings, see "Configure Context Hub Data Source Settings" in the *Context Hub Configuration Guide*.
Time Window	The time window based on the value set for the Query Last field in the Configure Data Source Settings dialog. By default, the alert data for last 7 days is fetched.
Last Updated	The time when contextual data was last fetched from data source.

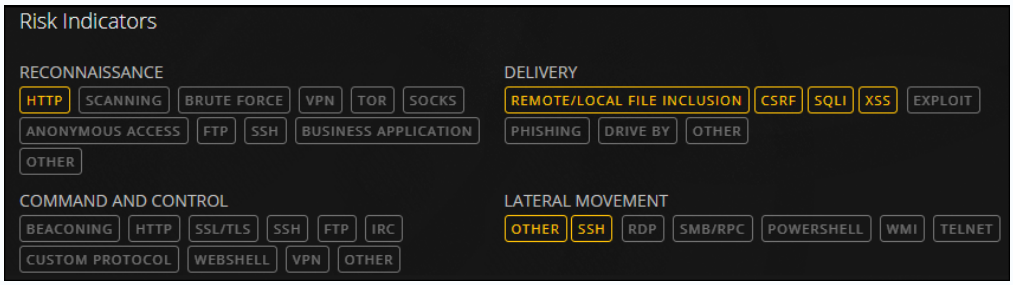
Live Connect Tab

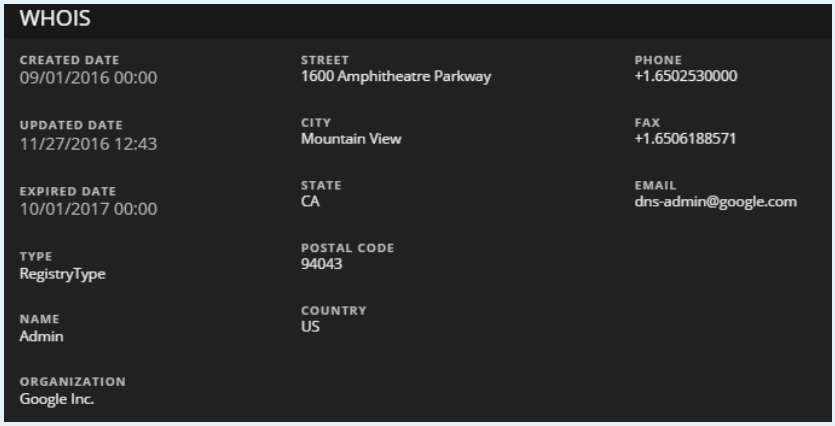
The following figure is an example of a Context Panel for Live Connect, and the table describes the information displayed.

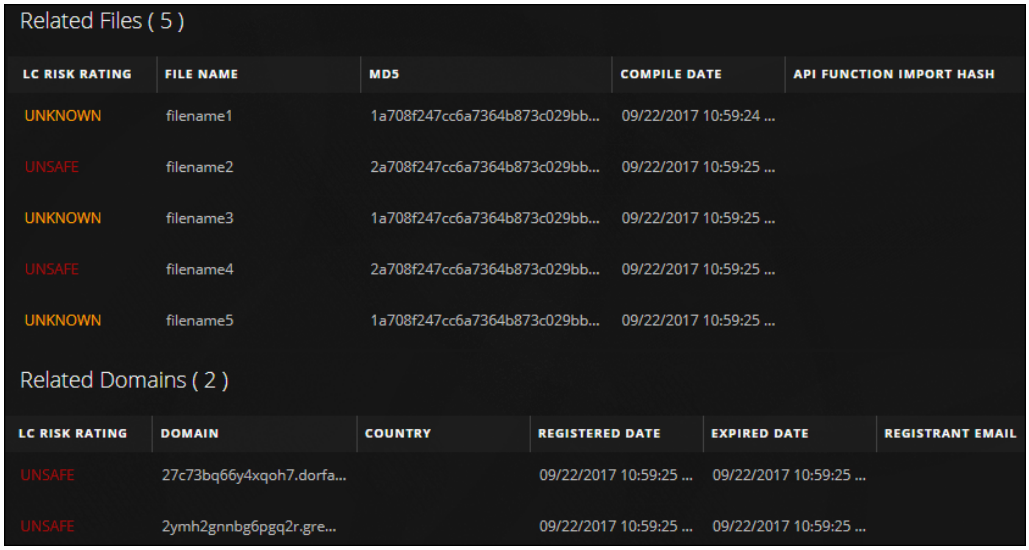


Field	Description
Review Status	The review status of the selected Live Connect entity (IP, file, or domain) based on the analyst activity. This gives the visibility of the analyst activity within an organization. **Status** Below are the types of status: • **New:** Lookup results for an IP address are viewed for the first time within the organization. • **Viewed:** Any analyst within the organization has already viewed the lookup results for an IP address. • **Marked as Safe:** Any analyst within the organization has already viewed the lookup results and marked the IP address as safe. • **Marked as Risky:** Any analyst within the organization has already viewed the lookup results and marked the IP address as risky.
Risk Assessment	The risk assessment for the selected Live Connect entity (IP, file, or domain) based on the Live Connect analysis and analyst feedback. The Risk Assessment categories are: • **Safe:** The Live Connect entity is considered to be safe. • **Unknown:** Live Connect does not have enough information about this entity to calculate the risk. • **High Risk**: Marked as high risk based on the analysis and risk reasons provided by the community. Entities marked as high risk require immediate attention. • **Suspicious**: Marked as suspicious based on the analysis and risk reasons provided by the community. The analysis indicates potentially threatening activity that requires action. • **Unsafe:** Marked as unsafe based on the analysis and risk reasons provided by the community. The entity is rated as High Risk, Suspicious, or Unsafe and displays the associated risk reasons accordingly.

Field	Description
Risk Assessment Feedback	 Risk Assessment Feedback allows the analyst to submit threat intelligence feedback about an entity to the Live Connect server. • **Analyst Skill Level** Below are the Analyst skill level options: ◦ **Tier 1** - Analysts at this level define procedures for remediation, and decide if an incident should be escalated to other areas in a Security Operation center (SOC). This is the default value. ◦ **Tier 2** - Analysts who investigate incidents and capture intelligence from an investigation to feed back into the various workflows in a SOC. ◦ **Tier 3** - Analysts who share the investigation results to the SOC organization. They generally manage incidents and have a wide breadth and depth of skills and tools necessary for incident response. **Note:** While creating a new user for NetWitness (Analyst), an administrator should be able to identify the user as Tier 1, Tier 2, or Tier 3 Analyst. • **Risk Confirmation** - The risk confirmation for the selected Live Connect entity (IP, file, or domain). The Risk confirmation categories are: ◦ **Safe**: The Live Connect entity is considered to be safe. ◦ **Unknown**: The analyst does not have enough information to provide a risk confirmation ◦ **High Risk**: Marked as high risk based on the analysis and risk reasons provided by the community. Entities marked as high risk require immediate attention. ◦ **Suspicious**: Marked as suspicious based on the analysis and risk reasons provided by the community. The analysis indicates potentially threatening activity that requires action. ◦ **Unsafe**: Marked as unsafe based on the analysis and risk reasons provided by the community. • **Confidence Level** - The confidence level of an analyst in providing feedback for the Live Connect entity. The confidence level categories are: High, Medium, and Low. • **Risk Indicator Tags** - Allows you to select a tag category based on the analysis.

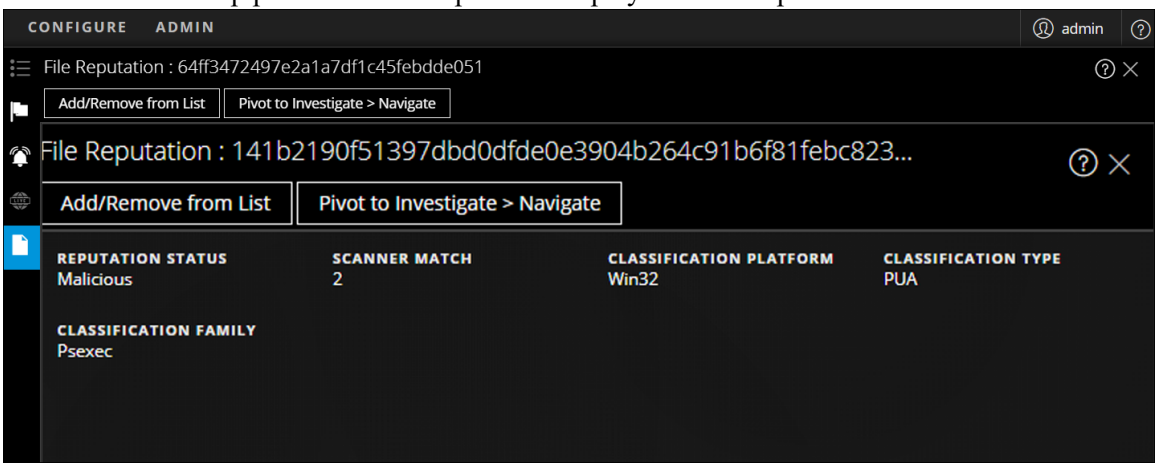
Field	Description
Community Activity	Community activities such as: • Date first seen in the community. • Time since the IP/File/Domain was seen for the first time (Current time - First seen time). **Trending Community Activity:** If the IP address is known within the RSA community, a graphical representation of the community activity trend is displayed for the following: • Users (in %) who have viewed the IP address in the Live Connect community over time. • Users (in %) who submitted feedback for the IP address. • Users (in %) who marked the IP address as unsafe over time.
Risk Indicators	 Risk indicators are highlighted based on the tags that are assigned by the community to the entities (IPs, Files, or Domains). The tags are categorized as follows: Reconnaissance, Delivery, Command and Control, Lateral Movement, Privilege Escalation, and Packaging and Exfiltration. These tags are samples and vary based on the inputs received from the community on the Live Connect server. The analyst can choose the appropriate risk indicator tags while providing the review feedback. A highlighted tag indicates that the selected entity is associated with that particular category and tag. Clicking a highlighted tag displays the description of the tag.
Identity	Provides the following identity information for the selected entity or meta value: For IP address: Autonomous System Number (ASN), Prefix, Country Code and Country Name, Registrant (Organization), and Date. For File Hash: File Name, File Size, MD5, SH1, SH256, Compile Time, and Mime Type. For Domain: Domain Name and Associated IP Address.
Certificate Information	Provides the following certificate information for the selected file hash: Certificate Issuer, Validity of the Certificate, Signature Algorithm, and Certificate Serial Number.

Field	Description
WHO IS Information	 The **WHO IS** information provides the ownership details for a given domain. The following information about the domain owner is displayed: Created Date, Updated Date, Expired Date, Type (Registration Type), Name, Organization, Address with Postal code, Country, Phone, Fax, and Email.
Related Files	Related Files are displayed for entity types IP and Domain. A list of known associated files is displayed along with the following information: Live Connect Risk Rating (Safe, Risky, or Unknown), File Name, MD5, Compile Time and Date, API Function, Import Hash, and Mime Type.
Related Domains	Related Domains are displayed for entity types IP and Files. A list of known associated domains is displayed along with the following information: Live Connect Risk Rating (Safe, Risky, or Unknown), Domain Name, Country Name, Registered Date, Expired Date, and Registrant Email address.

Field	Description
Related IPs	 Related IPs are displayed for entity types Domain and Files. A list of known associated IPs is displayed along with the following information: Live Connect Risk Rating (Safe, Risky, or Unknown), IP Address, Domain Name, Country Code and Country Name, Country Name, Registered Date, Expired Date, and Registrant Email address.

File Reputation Tab

The Context Lookup panel for File Reputation displays the file reputation status of a file.

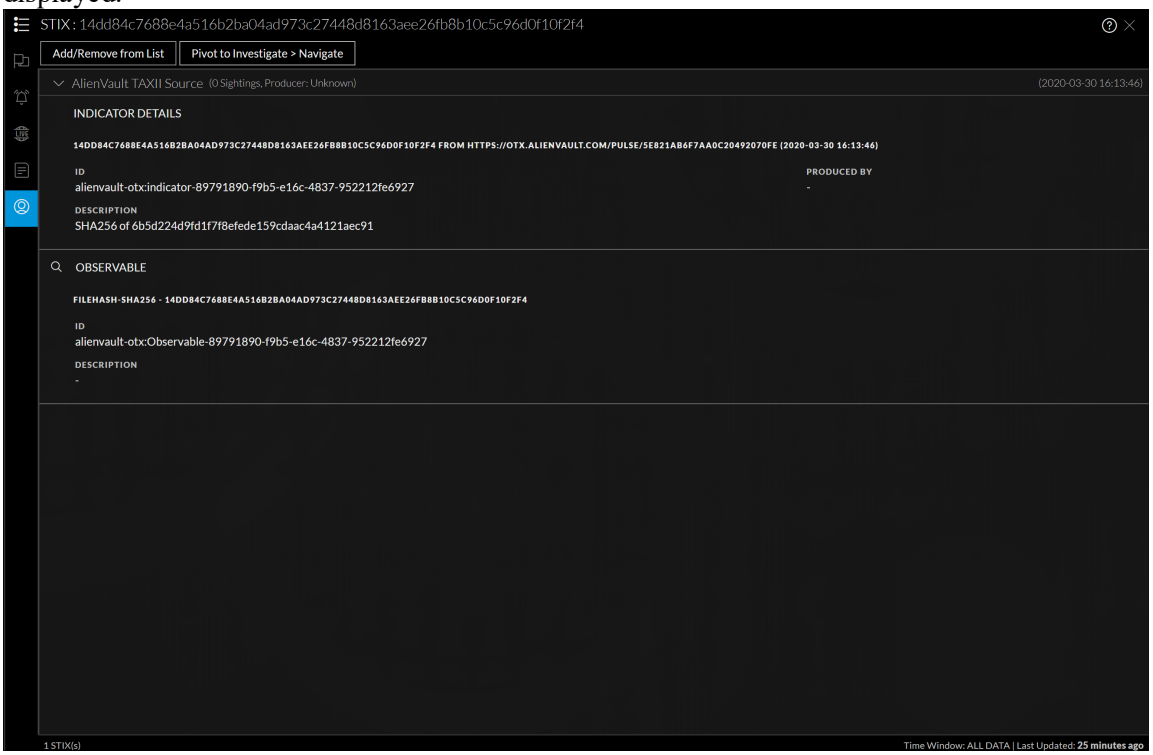


Field	Description
Reputation Status	Reputation Status of filehash. For more information about reputation status, see "View Reputation of files" in the *UEBA User Guide*.
Scanner Match	Number of scanners that detected malware or suspicious activity in the last scan.

Field	Description
Classification Platform	Classification for the queried filehash based on the platform. For example, the platform can be Win 32.
Classification Type	Classification for the queried filehash based on the type.
Classification Family	Classification for the queried filehash based on the malware family name.

TI Tab

The following figure is an example of a Context Panel for TI, and the table describes the information displayed.



Field	Description
Data Source name	Displays the STIX data source name from where the data is retrieved.
Timestamp	The time when the event was created.
Indicator Details	Indicator Title: Displays the details that contains a pattern that can be used to detect suspicious or malicious cyber activity. ID: Displays the ID of the selected indicator. Produced by: Displays the user role who requested for the STIX data. Description: Displays details about the selected IP address which are being watch listed.

Field	Description
Observable	Observable Title: Displays and conveys information about cyber security related entities such as files, systems, and networks using the STIX Cyber-observable Objects (SCOs). ID: Displays the ID of the selected observable.
(Optional) SightingsREST	Sightings Title: Displays the name of the sighting source. Confidence: Displays the criticality of the sighting. Reference: Displays the reference URL of the sighting source.

REST API Tab

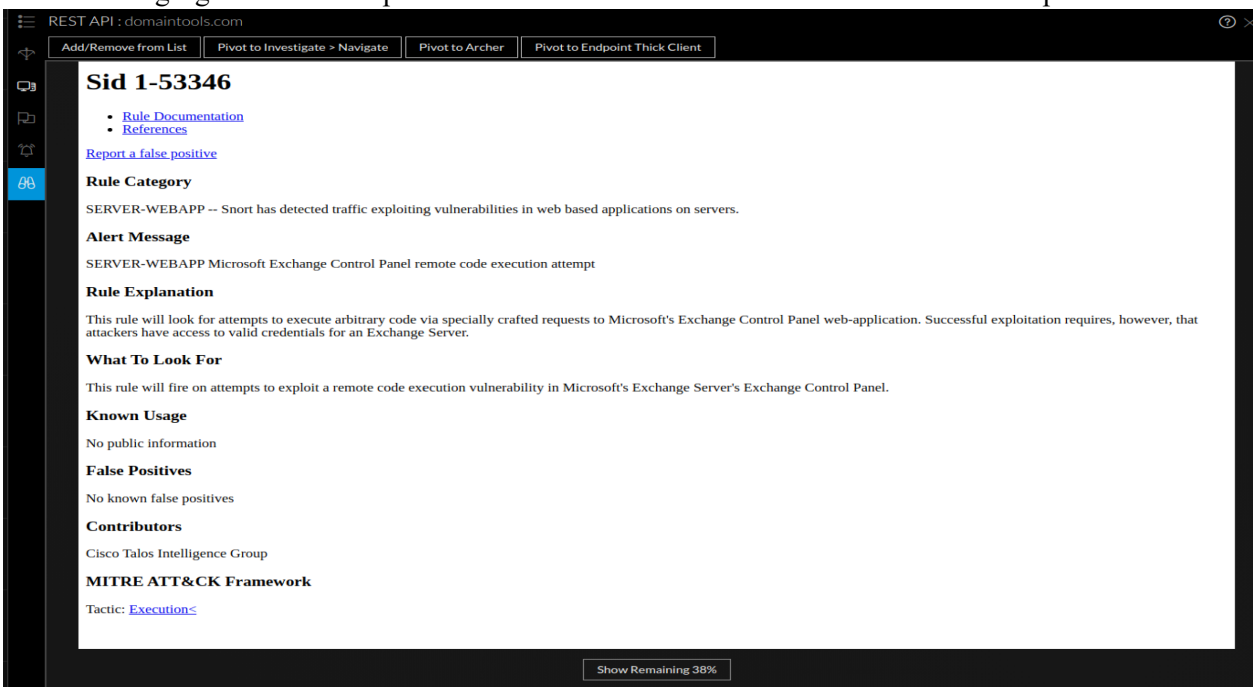
The Context Lookup panel for REST API shows HTML or JSON response (based on the response type configured) associated with the selected entity or meta value.

Note: For JSON response type, the fields that are mapped with friendly names (during REST API configuration) are only displayed for context Lookup. If you have not mapped any fields, all fields are displayed for context lookup.

The following figure is an example of the Context Panel for REST API with JSON response:



The following figure is an example of the Context Panel for REST API with HTML response:

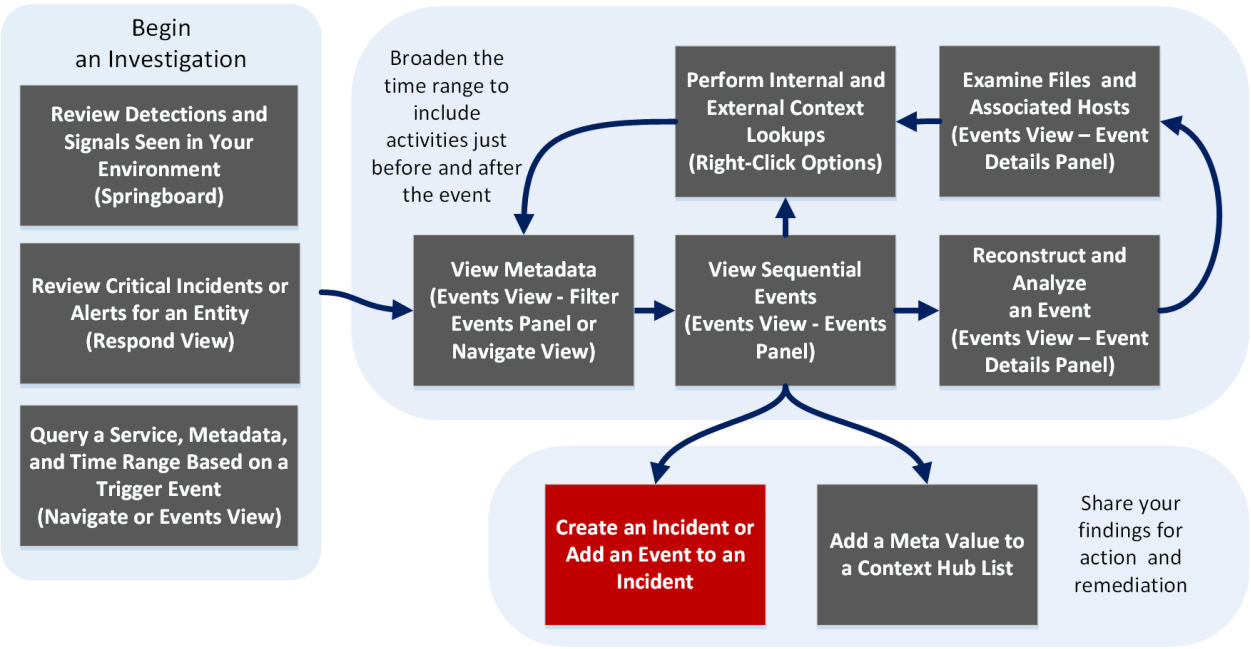


Create an Incident Dialog

In the Create an Incident dialog, analysts can create an incident from selected events in the Events view. The incident is then available to incident responders working in Respond.

To access this dialog, while investigating a service in the Investigation > Events view, select **Incidents > Create New Incident** from the toolbar.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)

User Role	I want to ...	Show me how
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident*	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

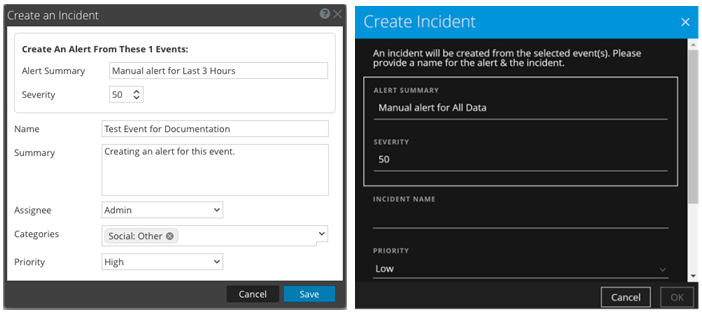
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Navigate View
- Legacy Events View

Quick Look

The following figure is an example of the Create an Incident Dialog, and the features are described in the table.



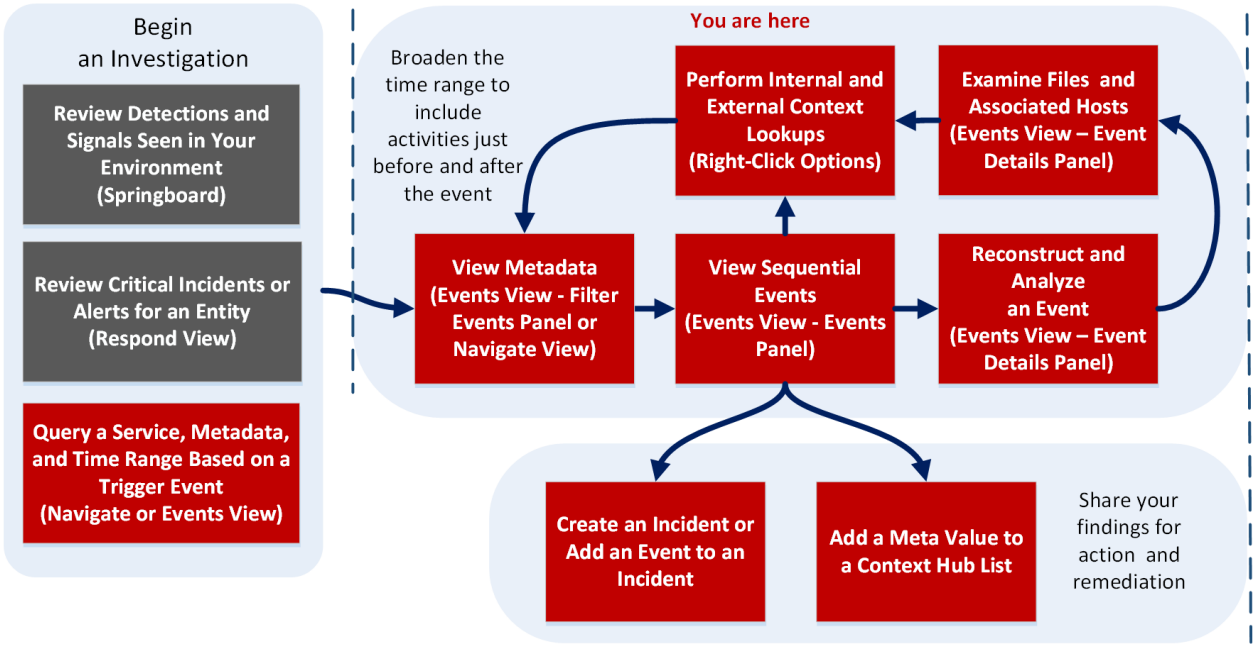
Feature	Description
Create Summary from These Events	The Alert Summary field is filled by the query that produced the select alerts, which you selected to create this incident. The Severity field reflects the Severity of the selected alert, an integer between 1 and 100.
Name	(Required) Specifies a name to identify the incident. In the example, the name is Sample Incident. You can provide a name that clearly identifies the nature of events that will be added to this incident
Summary	(Optional) Specifies a description for the incident. A good summary clearly identifies the incident for other analysts and responders.
Assignee	(Optional) Assigns the incident to a user in the SOC. Clicking Assignee opens a drop-down list showing the user names of SOC personnel who respond to incidents.
Categories	(Optional) Identifies categories of incidents. Clicking Categories, opens a drop-down list of Incident categories and subcategories. You can select one or more categories to which the incident belongs. Categories fall into these major groups: Environmental, Error, Hacking, Malware, Misuse, and Social.
Priority	Identifies the priority for the incident. Clicking Priority opens a drop-down list of priorities: Critical, High, Medium, or Low displayed in the drop-down list.
Cancel	Closes the dialog without saving changes.
Save	Saves the incident and closes the dialog. A message confirms that the incident was created successfully.

Events View

In the Events view analysts can view a sequential list of network, log and endpoint events, select an event for reconstruction and analysis, and view the raw event and metadata with interactive features that enhance the ability to see meaningful patterns in the data. In Version 11.5 and later, you can drill into metadata for the listed events. The Events view offers packet, file, host, text, log, and email reconstruction. When you open a web reconstruction of an event, the same web reconstruction used in the Legacy Events view is displayed.

Workflow

The following figure is a high-level workflow illustrating the tasks you can do in NetWitness Investigate, with the Events view tasks highlighted in red.



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*

User Role	I want to ...	Show me how
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata*	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)
Threat Hunter	view sequential events*	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident*	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

Quick Look

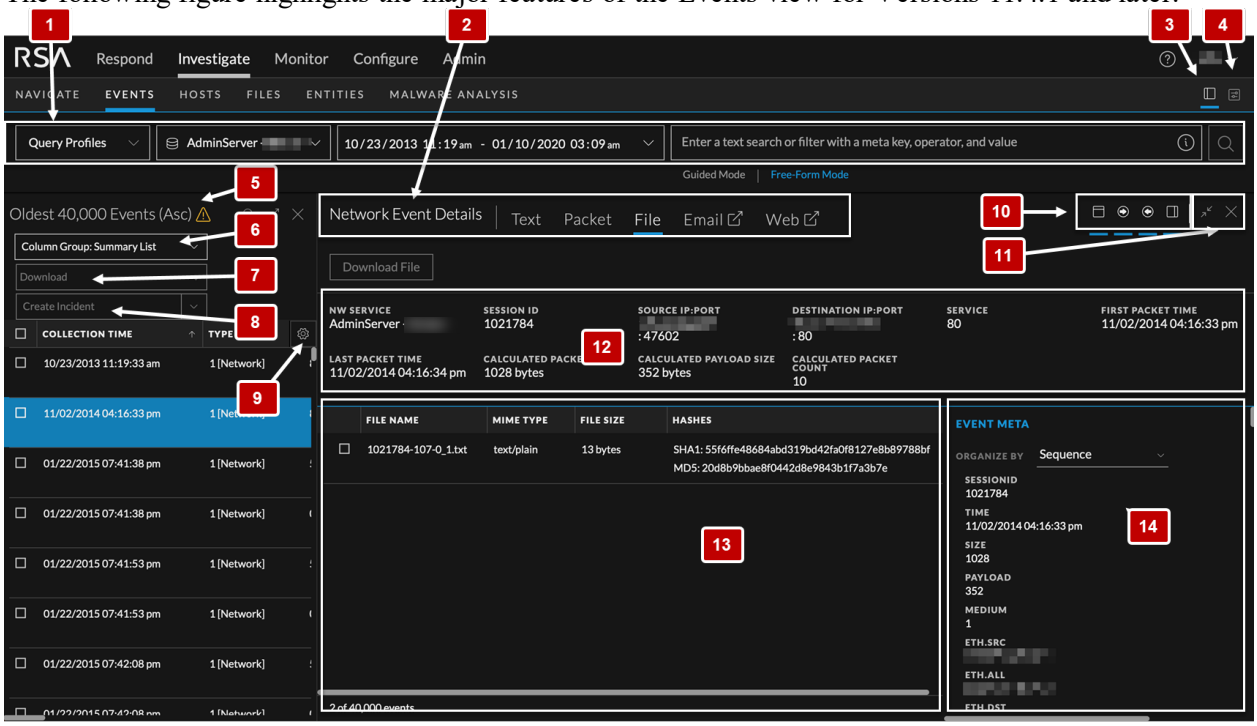
There are multiple access points to this view, which are described in Begin an Investigation in the Events View. If you access the Events view from the Respond view, you can see the analysis for a selected event in an incident. The options are a subset of the options available when you open an event from within the Investigate view. To get complete functionality and examine other events, you can go to the Event view directly (INVESTIGATE > Event).

The Events view lists events in ascending order by time in the Events panel. The events displayed can be results for the drill point in the Navigate view or Legacy Events view, or results for a query entered in the Events view query bar.

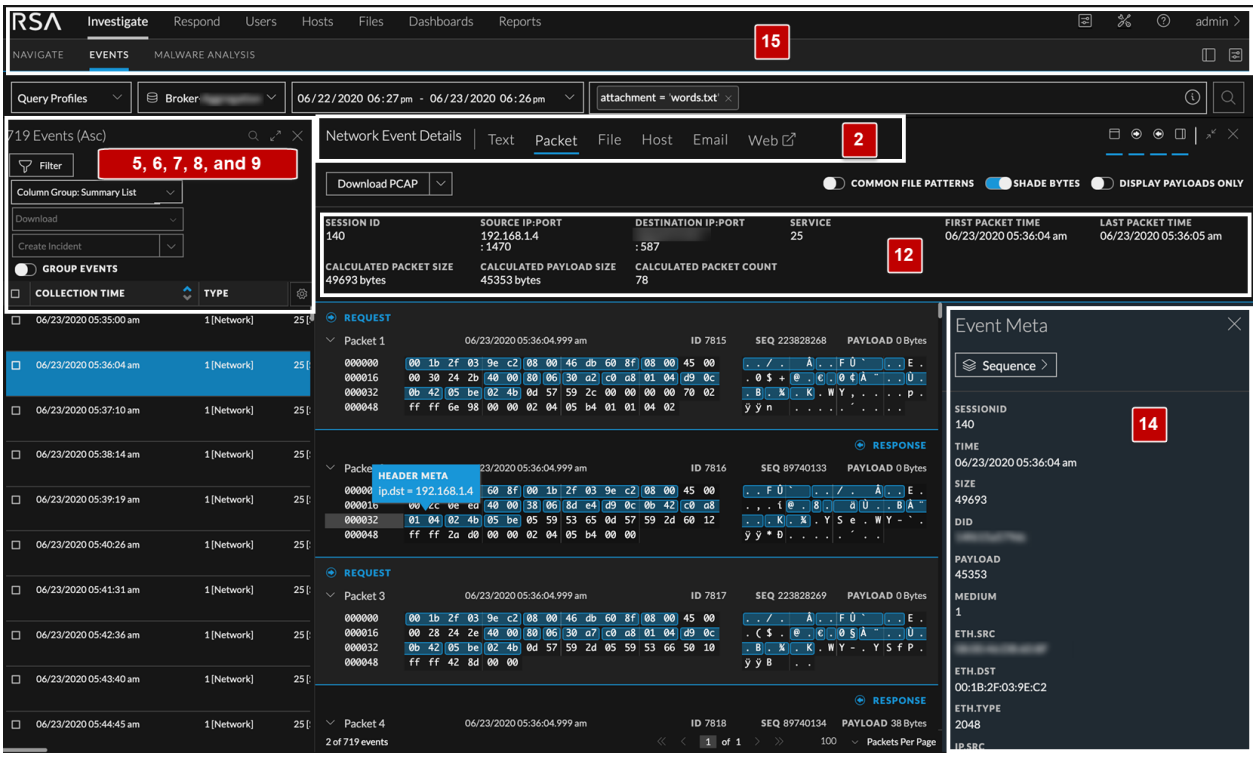
Input fields for a query are displayed so that you can select a service and time range, and type an optional query. When you submit a query, the service being investigated counts the results up to a limit of 10,000 events, and 10,000 network, log, and endpoint events are loaded in the Events panel. Different columns are displayed, depending on the selected column group. You can rearrange and resize the columns, choose a built-in or custom column group, and choose individual columns that you want to see. When you find an event of interest, clicking the event opens the reconstruction in a new panel (Packet, Text , or File).

Note: For versions earlier than 11.3, the first 100 events are loaded. You can scroll through the list and click **Show Next 100 Events** at the bottom of the list. If the next page contains fewer than 100 events, the button changes to reflect the number of remaining events.

The following figure highlights the major features of the Events view for Versions 11.4.1 and later.



This figure illustrates the changed elements in Version 11.5 and later.



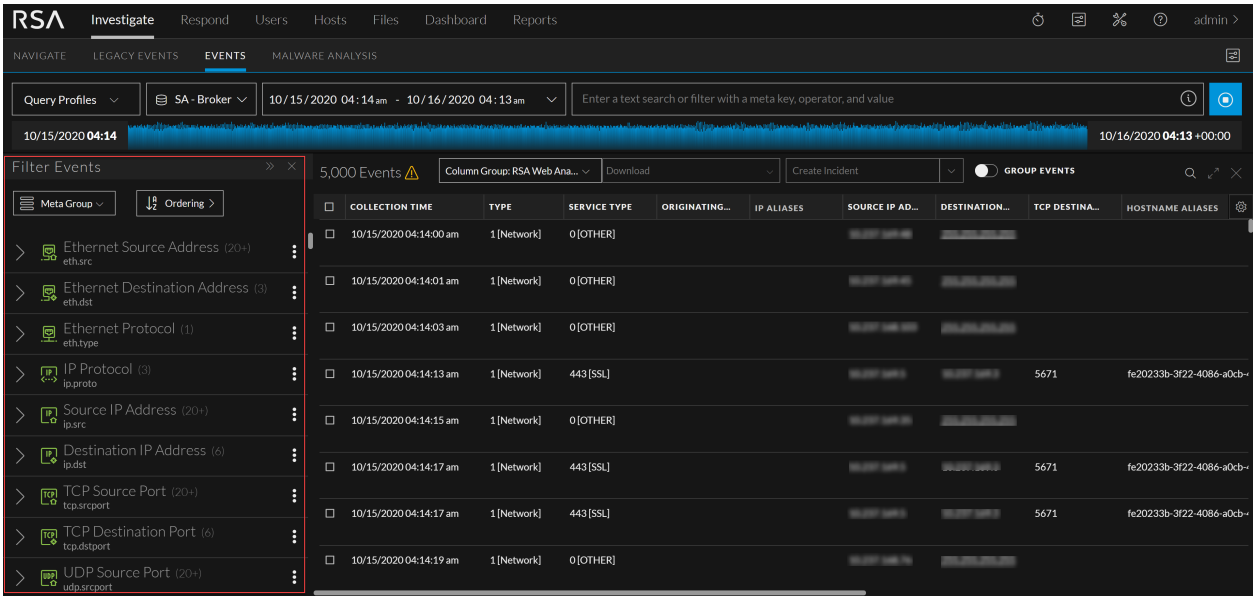
1	**Query Bar**: When a service is selected, displays the service selector, time range selector, and the queries you have entered. You can select a service as described in Begin an Investigation in the Events View and refine the query as described in Filter Results in the Events View. Clicking submits the query and sends a request to the selected service to load the data. In Version 11.3 and later, clicking the (console icon) opens the query console, where detailed status of the query is provided (see Events View below).
2	The type of event being analyzed and the type of reconstruction are reflected in the heading. • These are the event types: **Network Event Details**, **Log Event Details**, or **Endpoint Event Details**. • The types of analysis available for the event type are Text, Packet, File, Host, Email, and Web. Network events can use all types of analysis: text, packet, file, and email (Version 11.4.1 and later). Log and endpoint events use only text analysis. The email (Version 11.4.0.x and earlier) type and web type open the current event as an email or web reconstruction in the Events view. For details, see Examine Event Details in the Events View.
3	Reopens the Events panel if you have closed it. For details, see Analyze Events in the Events View.
4	Sets preferences for the Event view (see Configure the Events View).
5	The Events panel title. • In Version 11.3 and later, the Events panel title is slightly different than the title in prior versions, and a row number indicator has been added. The title lists the number of events and sort order; for example, **24,000 Events (Asc)** means that 24,000 events

	were found and they are listed in ascending order by time. If more than 10,000 events are found, only the oldest 10,000 events are displayed in ascending order, and an amber triangle highlights the fact that not all events were loaded. This may indicate that you need to refine the query. For more information about refining the events listed here, see Filter Results in the Events View. • Versions prior to 11.3 simply list the number of events found, and you can load 100 of them at a time. In Version 11.4 and later, clicking  opens the Find Text in Table dialog.
6	The Column Group drop-down lists built-in and custom column groups that you can apply to the Events panel. Built-in column groups are sometimes updated between one version and the next. Some examples of built-in column groups are Email, Endpoint Analysis, Malware Analysis, Outbound HTTP, Outbound SSL/TLS, and Summary List. Summary List is the default column group. For details, see Use Columns and Column Groups in the Events List.
7	The Download drop-down menu lists the available options for downloading event data. The options are Log, Visible Meta, and Network (see Downloading and Acting Upon Results. You can change the preferred format of the event type data in the Event Preferences dialog (see Configure the Events View).
8	The Create Incident button enables you to create incidents from events. The Add to Incident button enables you to add selected events to an open and existing incident (see Add Events to an Incident in the Events View and Add Events to an Incident in the Legacy Events View).
9	Displays the column selection settings to select the individual columns displayed in the Events panel. For details, see Use Columns and Column Groups in the Events List.
10	Controls to show or hide the Event Header, show or hide requests and responses, and open the Event Meta panel. For details, see Analyze Events in the Events View.
11	Controls to change the size of the panel and close the panel. For details, see Analyze Events in the Events View.
12	The Event Header provides summary information about the event you are currently analyzing. The selected event is highlighted in the Events panel with a blue background. The summary information is different for the different event types (packet, log, and endpoint). In Version 11.5, the redundant NW Service is removed.
13	The event data for the event you are currently analyzing.
14	The Event Meta panel is redesigned in Version 11.5, but has the same functions as in Version 11.4. The Event Meta panel lists the meta keys and values found in the data. This data can be sorted in two ways - Alphabets or Sequence. Some metadata are searchable; they have a binoculars icon, which you can click to see the associated data highlighted in the event data (see Analyze Events in the Events View). • For a packet, the data is called a payload and is displayed in the form of a request and response. • For a log event, the data is a line of text from the raw log. • For an endpoint event, the event data is relevant to data from the NetWitness Endpoint agents running on hosts in the network. It may be a single process, driver, DLL, file (executable), service, or autorun, and information related to logged-in users. (See the *NetWitness Endpoint User Guide* for complete information about endpoint event data.)
15	The Version 11.5 main menu for RSA NetWitness Platform has relocated Hosts, Files, and Users (Entities) options for easier access.

Filter Events Panel

The Filter Events panel is a beta feature added from Version 11.5. Clicking the Filter button (Filter) in the Events panel, opens the panel to provide a view of meta keys and meta values found in the data set. (Version 11.6) By default, the Filter Events panel is open in the Events view. The user preference (open, closed, or fully expanded) is saved across sessions and logins. See Drill into Metadata in the Events View (BETA) for more information about drilling into metadata.

Note: (Version 11.6) By default, the Filter Events panel is open in the Events view. The last used state of the panel (narrow or fully expanded) is saved throughout the session and across logins. Also, the Filter Events panel provides additional contrast between meta keys, meta values, and meta counts to improve readability.



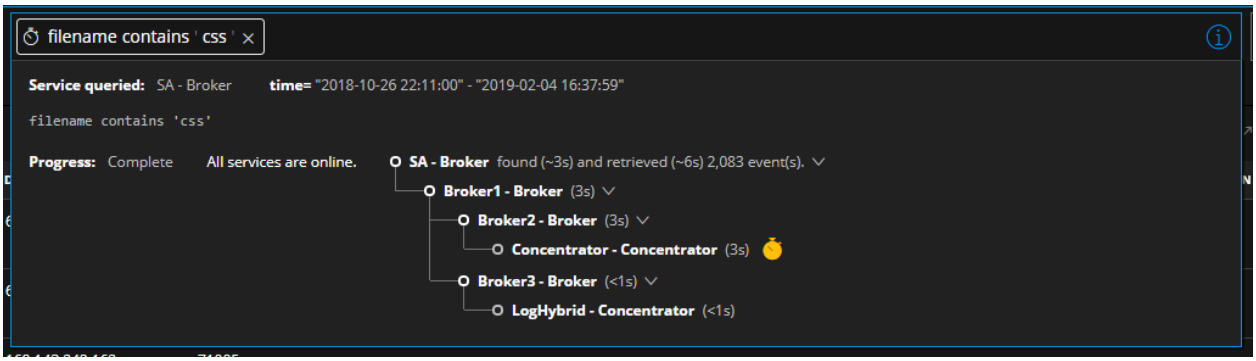
Meta Groups Menu

With the Filter Events panel open, you can select a meta group to define the meta keys displayed in the Filter Events panel.The Default Meta Keys meta group is in effect the first time you log in. If you selected a different meta group the last time you logged in, it remains in effect until browser cache is cleared. In Version 11.5.1, the meta group you selected previously is not stored in browser cache so it remains in effect until you change it. See Use Meta Groups to Focus on Relevant Meta Keys for details about meta groups.

Ordering Menu	With the Filter Events panel open, you can look at two parameters for each value: the event count or the event size. Each meta key entry includes either the event count or the event size in parentheses after the value. In both cases, there are four options for ordering: • By default, the meta keys are displayed using the Event Count > Descending by Total Count method. When showing the event count for each value, you can order by Descending by Total Count, Ascending by Total Count, Ascending by Value, and Descending by Value. • When you prefer to see the size of the event that contains the value, you can use one of the four Event Size ordering options: Descending by Total Size, Ascending by Total Size, Ascending by Value, and Descending by Value.
Meta Key options button (⋮)	The Meta Key options button offers actions that you can take on an individual meta key. In Version 11.5, the only action is to copy all of the visible meta values for a meta key.
Meta Key List	An icon before each meta key name identifies the indexing method for the key. The indexing method determines the types of interactions and queries possible using that meta key. • This meta key is indexed by value: Action Event [action] (40+). The green color indicates that the all available interactions and queries are supported. You can see the available interactions in the context menu by right-clicking the meta value. • This meta key is indexed by meta key: Bytes Sent [bytes.src]. The yellow color is a clue that a subset of available interactions is supported, and queries on this meta key may take longer than meta keys that are indexed by value. You can see the available interactions in the context menu by right-clicking the meta value. • This meta key is not indexed: MAC Alias Record [alias.mac]. Values for non-indexed meta keys cannot be used to query. If you want to query a meta key that is not indexed, your administrator needs to edit the index file for the service to index the meta key by value or meta key.

Query Console

Clicking (the console icon) opens the query console, where detailed status of the query is provided.



In the query console, you can see which service, time range, and metadata was queried as well as real-time information about the status of the query and the services being queried. A progress bar indicates the query's completion percentage at the bottom of the console. The statuses let you know details about what is happening; for example, you can tell when the query is executing, queued, reading the index file for the queried service, retrieving events, and complete. All statuses and non-fatal messages are displayed as they come in, and the border color changes if a non-fatal error occurs. View Status of a Query provides additional details on this subject.

Several messages that may be displayed in the query console require additional explanation.

Message: Maximum value limit (`valueMax`) of %1% reached on meta key %2% in index slice %3%

Explanation: The `valueMax` property on the specified meta key has been reached in the index being queried. An administrator configures this inside the index files available in `ADMIN > Services > [Service Name] > Files > index-[service type].xml` or `index-[service type]-custom.xml`. As an example, the statement below from the index file states the meta key called `client` has a limit of 250,000 values by default.

```
<key description="Client Application" level="IndexValues" name="client"
format="Text" valueMax="250000" />
```

Message: The query on channel %2% was auto-canceled by the system for exceeding time usage limits. Check timeout values.

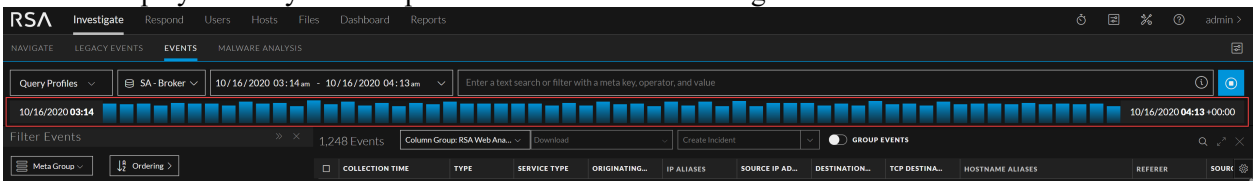
The server has a per-operation limit on execution time, and the requested operation exceeded the limit. To avoid this error, split the operation into smaller pieces, such as smaller time ranges.

Memory limit of %1% reached, controlled by setting max.query.memory

The server has a per-operation limit on memory utilization, and the requested operation exceeded the limit. The limit is related to the amount of memory in the server, which an administrator can adjust in **ADMIN > Services > [Service Name] > sdk > config**. To avoid this error, split the operation into smaller pieces, such as smaller time ranges.

Timeline

The timeline visualizes the events count that occurs at a specific instance. The timeline provides event counts so that you can see if the number of events increases drastically at a given point in time. The timeline displays activity for the specified service and time range as a bar chart.



The following are some key features of the timeline:

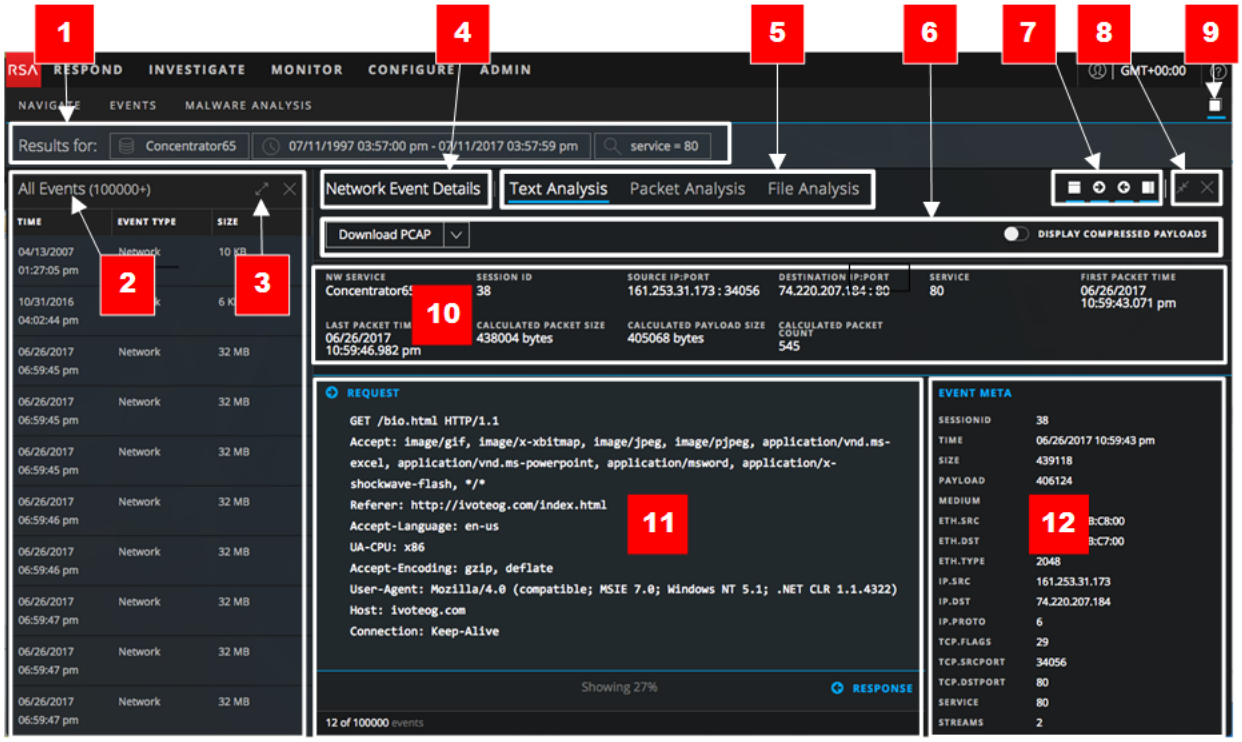
- The start time and end time are displayed on either side of the timeline to depict the matching events over the duration of the query. The timezone is also displayed on the right side of the timeline by showing the offset from UTC.
- The section of the timeline that matches the event data is highlighted. The data is highlighted up to the specified threshold. If the data is not sorted using time, the entire timeline is highlighted. Also, if the

data is not sorted in a specific order, it might highlight the individual bars in the timeline as events are not ordered using time.

- The query results that exceed the threshold are displayed in grey color. If events are sorted using time, the threshold highlight shifts from left (ascending) to right (descending) depending on your sort preference.
- Hover over the timeline to display the total number of events that occurred in the queried time duration. You can hover over the individual bars in the timeline to get the total number of events that occurred at a specific time.
- When you sort the events, the timeline does not refresh. You must run the query again to refresh the timeline.
- All capabilities of the timeline work when user is querying using the collection time (`time`) or when the preference is set to use the event time (`event.time`).

Quick Look for Version 11.0.0.x (End-Of-Life)

The following figure highlights the major features of the Event Analysis view for Version 11.0.0.x.



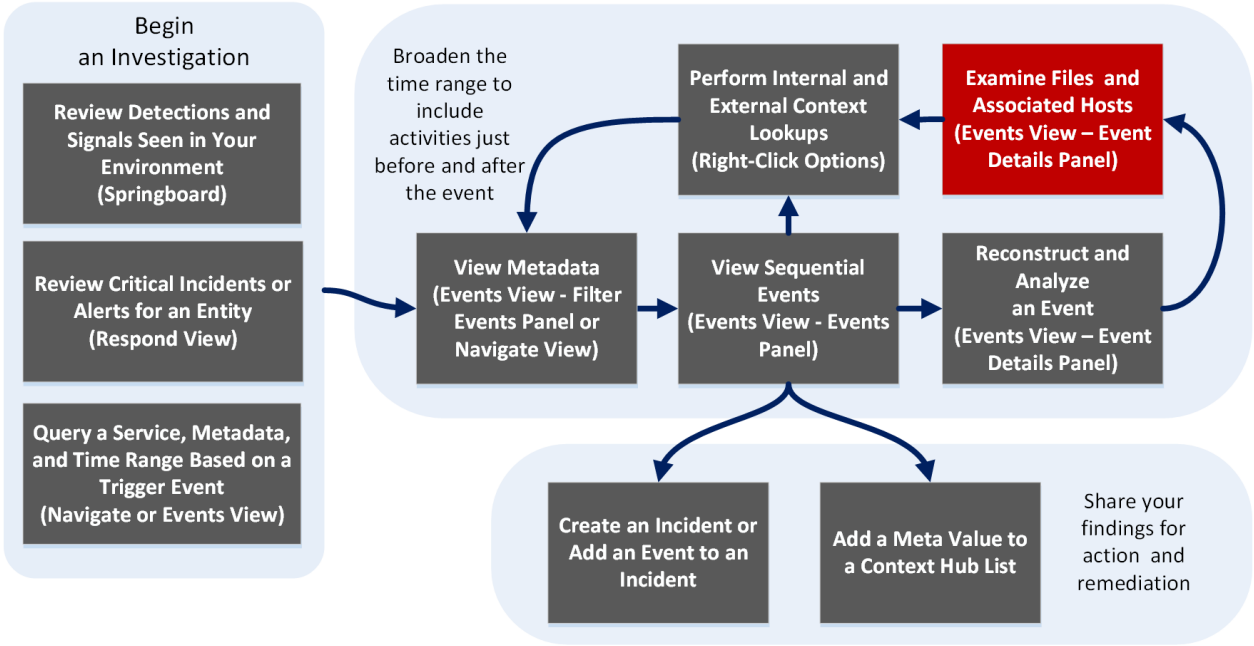
1	The read-only breadcrumb displays the selected service, time range, and query entered in the Navigate view or Events view.
2	This is a read-only list of events based on the query made in the Navigate or Events view. The Events panel includes a count of the events. You can rearrange and resize columns. You can scroll to the bottom of the list, and load more events (see Analyze Events in the Events View).
3, 8	Controls to change the size of the panel and close the panel.
4	The type of event being analyzed is reflected in the heading: Network Event Details, Log Event Details, or Endpoint Event Details. Each view is discussed in detail in Analyze Events in the Events View.

5	The types of analysis available for the event type. Network events can use all three types of analysis: text, packet, and file. Log and endpoint events use only text analysis.
6	These options vary for the different types of analysis. They are discussed in detail in Analyze Events in the Events View.
7	Controls to show or hide the Event Header, show or hide requests and responses, and open the Event Meta panel (12). These controls are described in Analyze Events in the Events View.
9	Reopens the Events panel or the Event Meta panel if you have closed it.
10	The Event Header provides summary information about the event. This information is different for the different event types (packet, log, and endpoint).
11	The event data (sometimes called a payload for packets). The event data for a log event or endpoint event is typically a line of text from the raw log rather than request and response shown for a packet.
12	The Event Meta panel lists the meta keys and values found in the data. Some meta data are searchable; they have a binoculars icon, which you can click to see the associated data highlighted in the event data (see Analyze Events in the Events View.

Events View - Email Tab

The Email tab is in the Event Details panel. Here you can view a list of email received and associated attachments for an event.

Workflow

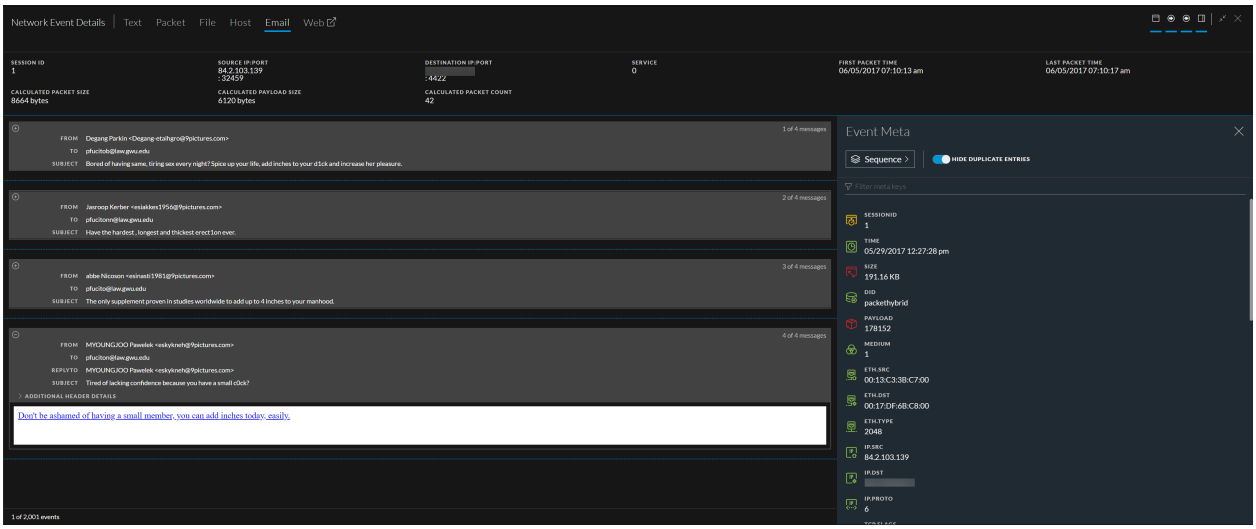


Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

Quick Look

The Email panel displays a list of emails associated with a network event. When an analyst opens the email, the email reconstruction is displayed along with the associated attachments and additional header details, if any.
The following figure is an example of an email reconstruction.



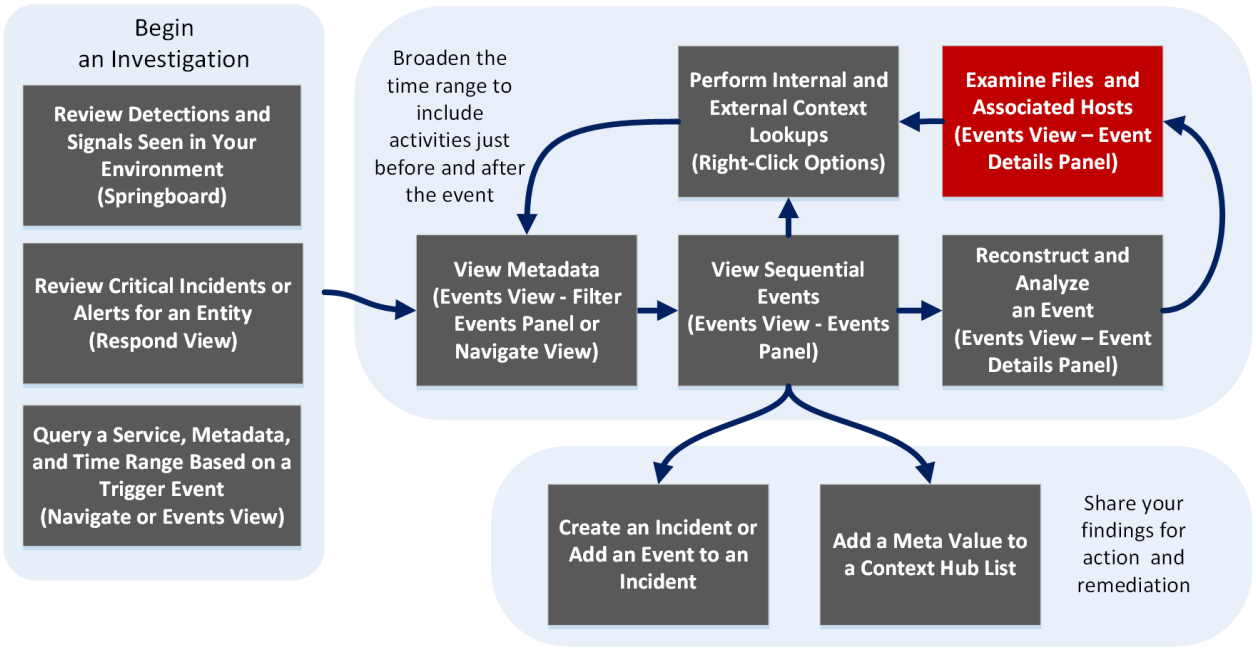
The following table describes all the fields within an email.

Field	Description
From	Displays the email address of the sender of the email.
To	Displays the email addresses of the recipients of the email.
CC (Carbon Copy)	Displays email addresses of additional recipients of the email. The field is displayed only if the sent email has any value and the email addresses are visible to the recipient.
BCC	Displays email addresses of additional recipients privately. This field is displayed only if the sent email has any value and the email addresses are not visible to the recipient.
Reply to	Displays the address designated to receive replies, the sender address.
Subject	Displays the subject of the email.
Attachments	Displays any files shared by the sender that can be downloaded by the recipient. This field is displayed only if email contains attachments. See Download Data in the Events View for details about downloading email attachments.
Additional Header Details	Provides additional details of the email event such as Received, Sender, Message-ID and others.

Events View - File Tab

The File tab is in the Event Details panel. Here you can safely view a list of files and download one or more files in an event.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View

User Role	I want to ...	Show me how
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

*You can perform this task in the current view.

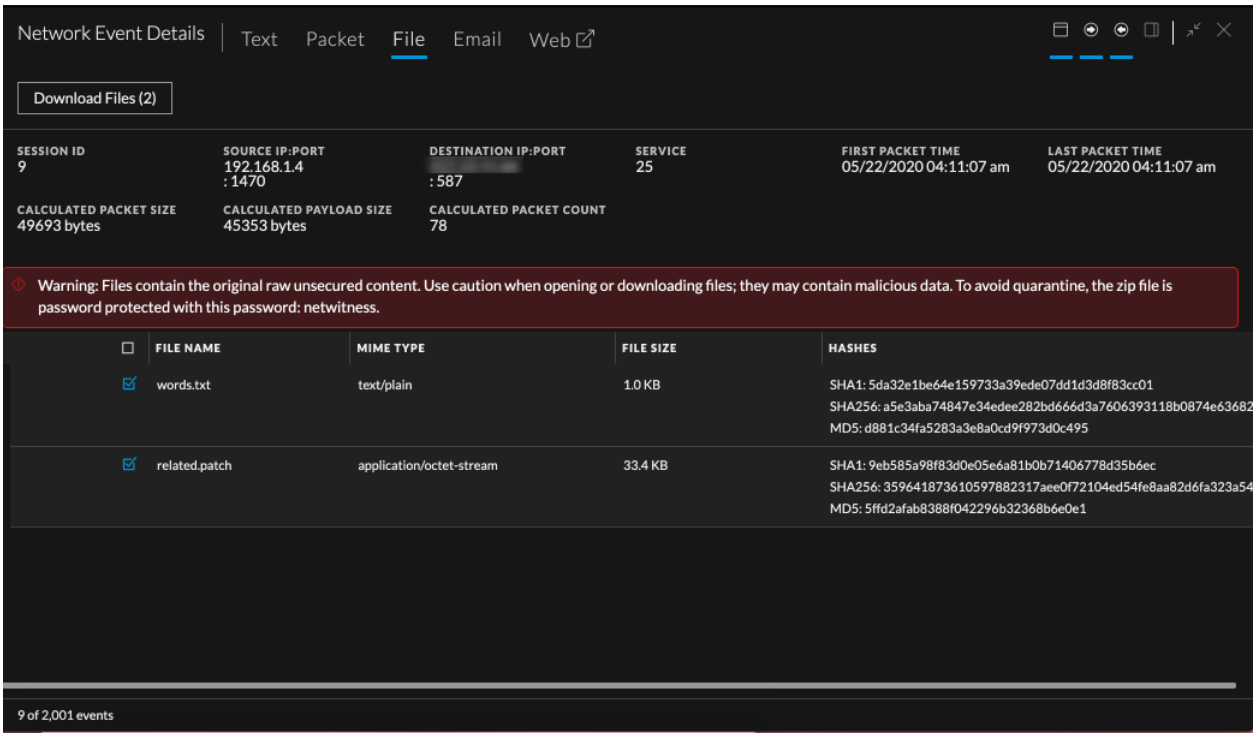
Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

Quick Look

The File panel displays a list of files associated with a network event. You can download files in this view.

Below is an example of the File panel.

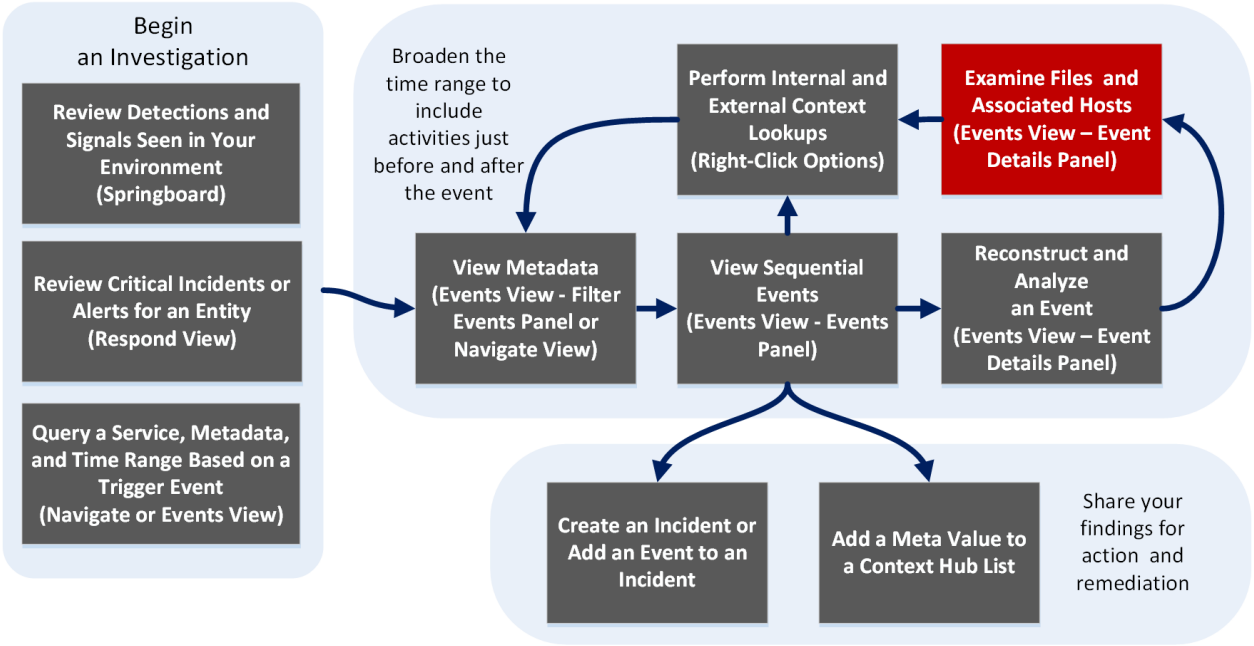


Feature	Description
Download Files button	Click to download one or more selected files.
Event Header	The Event Header displays summary information for the network event that contains the files.
Files List	Scrollable list of associated files that you can select and download.

Events View - Host Tab

The Host tab is in the Event Details panel. Here you can view network events enriched with endpoint data such as host and process triggered for the selected network event and other details such as risk score, reputation, and logged in user. The host panel is available for network events with endpoint data only.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)

User Role	I want to ...	Show me how
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

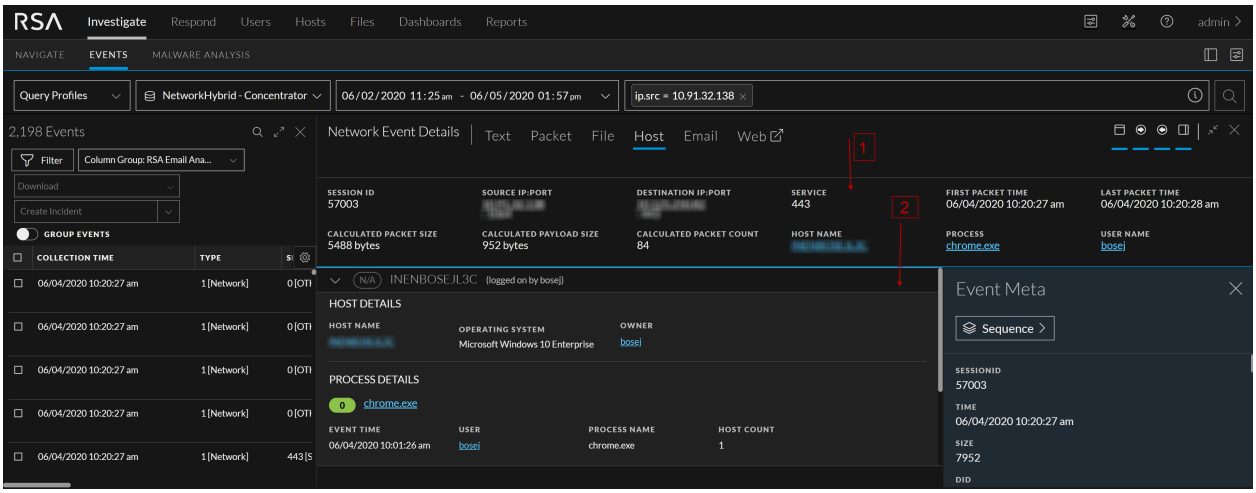
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

Quick Look

Below is an example of the Host panel with labeled features.



1 The event header displays the summary of network events enriched with endpoint data. It includes:

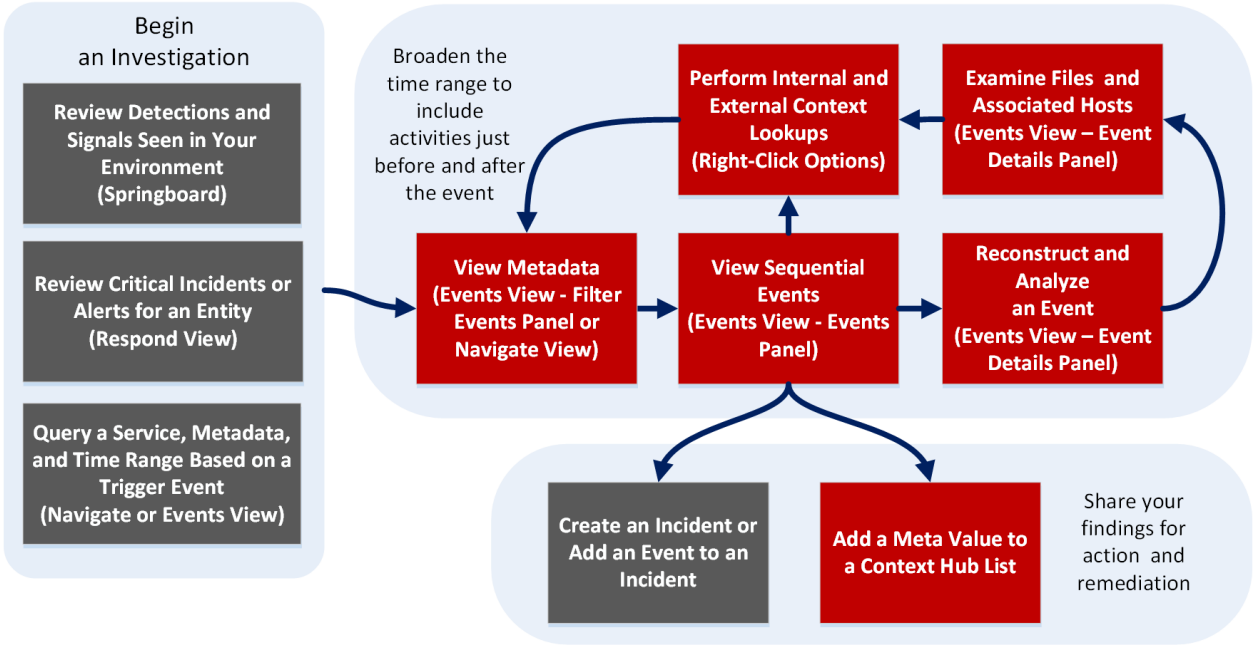
- Host - The host from where the event originated.
- Process - The source process which triggered the event.
- User - The user associated with the triggered process.

2 You can view additional details about the host and process. For more information, see Host Information.

Events View - Packet Tab

The Packet tab is in the Event Details panel. Here you can safely view and interactively analyze the packets and payload of an event.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata*	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)

User Role	I want to ...	Show me how
Threat Hunter	view sequential events*	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

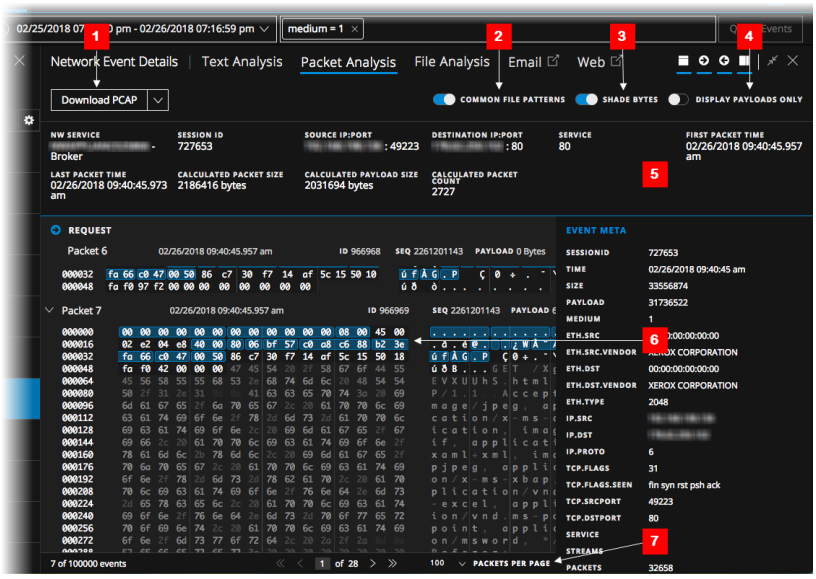
Quick Look

Only network events can be analyzed in the Packet panel. The Packet panel lists each packet in the event. The list of packets is scrollable. When you scroll, the packet or text identification information as well as the Request and Response labels remain visible rather than scrolling out of view.

In Version 11.1 and later, you can use pagination controls to go backward and forward through the pages, go to a specific page, and select the number of packets to display per page (50, 100, 300, or 500).

Each packet is displayed with shading and highlighting to help identify common file patterns: significant header and payload bytes, hexadecimal and ascii bytes, and common file signatures. In addition, you can adjust the request/response display, and display or hide the packet summary.

Below is an example of the Packet panel (formerly known as Packet Analysis) with labels to identify features. For details and examples of each feature, see Analyze Events in the Events View.

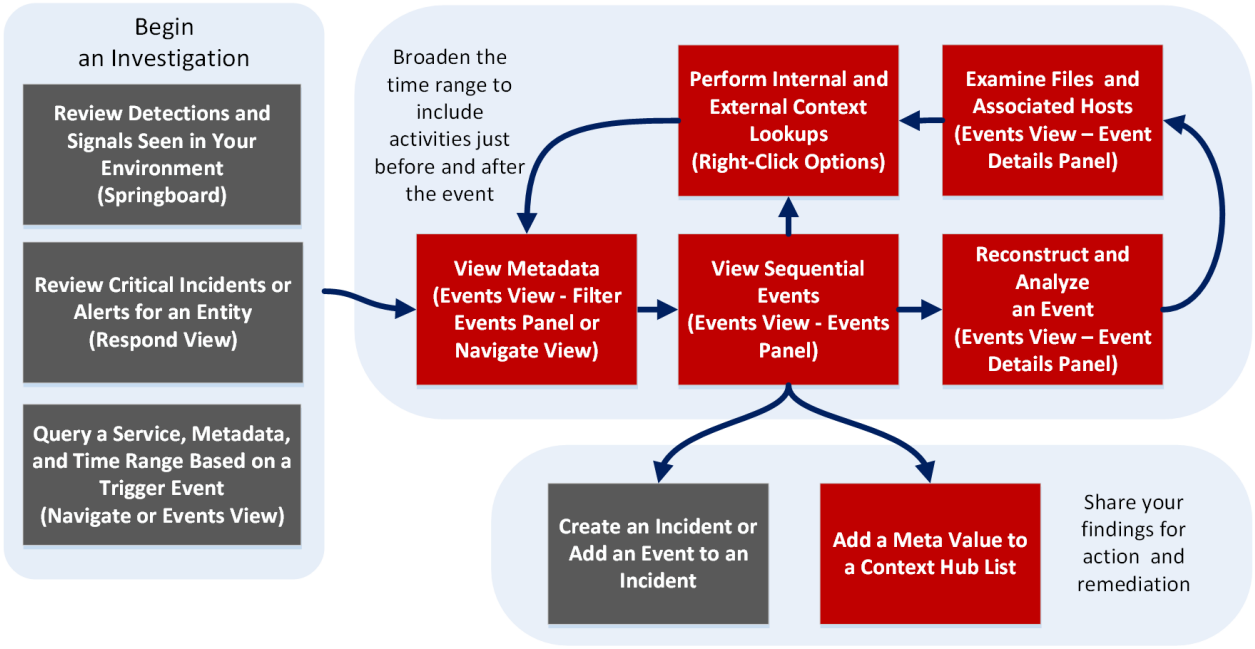


1	Options for exporting a network event. You can export a PCAP, all payloads, request payloads, or response payloads for deeper analysis and to share with others.
2	The option to identify common file signatures is activated by default. Common file signatures are highlighted in orange; hovering over the highlight reveals the file type.
3	The Shade Bytes option adds shading to identify the different hexadecimal bytes (00 to FF) using degrees of highlighting.
4	The option to display payloads only hides the packet headers, leaving more space for the payload.
5	The Event Header.
6	Significant bytes are highlighted in a blue background; as you move the cursor over the highlighting the meta data is displayed in a hover box.
7	(Version 11.1 and later) Packet pagination controls allow more flexibility in paging through a list of packets. When a control is unavailable, the image is dimmed; for example, when you are viewing page 1, the ⟪ and ⟨ controls are dimmed. ⟪ - Go to the first page ⟨ - Go to the previous page 1 of 206 - Go to a specific page ⟩ - Go to the next page ⟫ - Go to the last page  - Select the number of packets per page. If you are reconstructing large packets, lowering this limit can improve performance.

Events View - Text Tab

The Text tab is in the Event Details panel. Here you can safely view and analyze the raw text payload of an event. The Text reconstruction includes features that can show decompressed or compressed text, expand truncated entries, perform URL and Base64 encoding and decoding, and download network events, logs, and endpoint events. The text reconstruction is available for all types of events: network, log, and endpoint.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View

User Role	I want to ...	Show me how
Threat Hunter	view metadata*	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)
Threat Hunter	view sequential events*	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

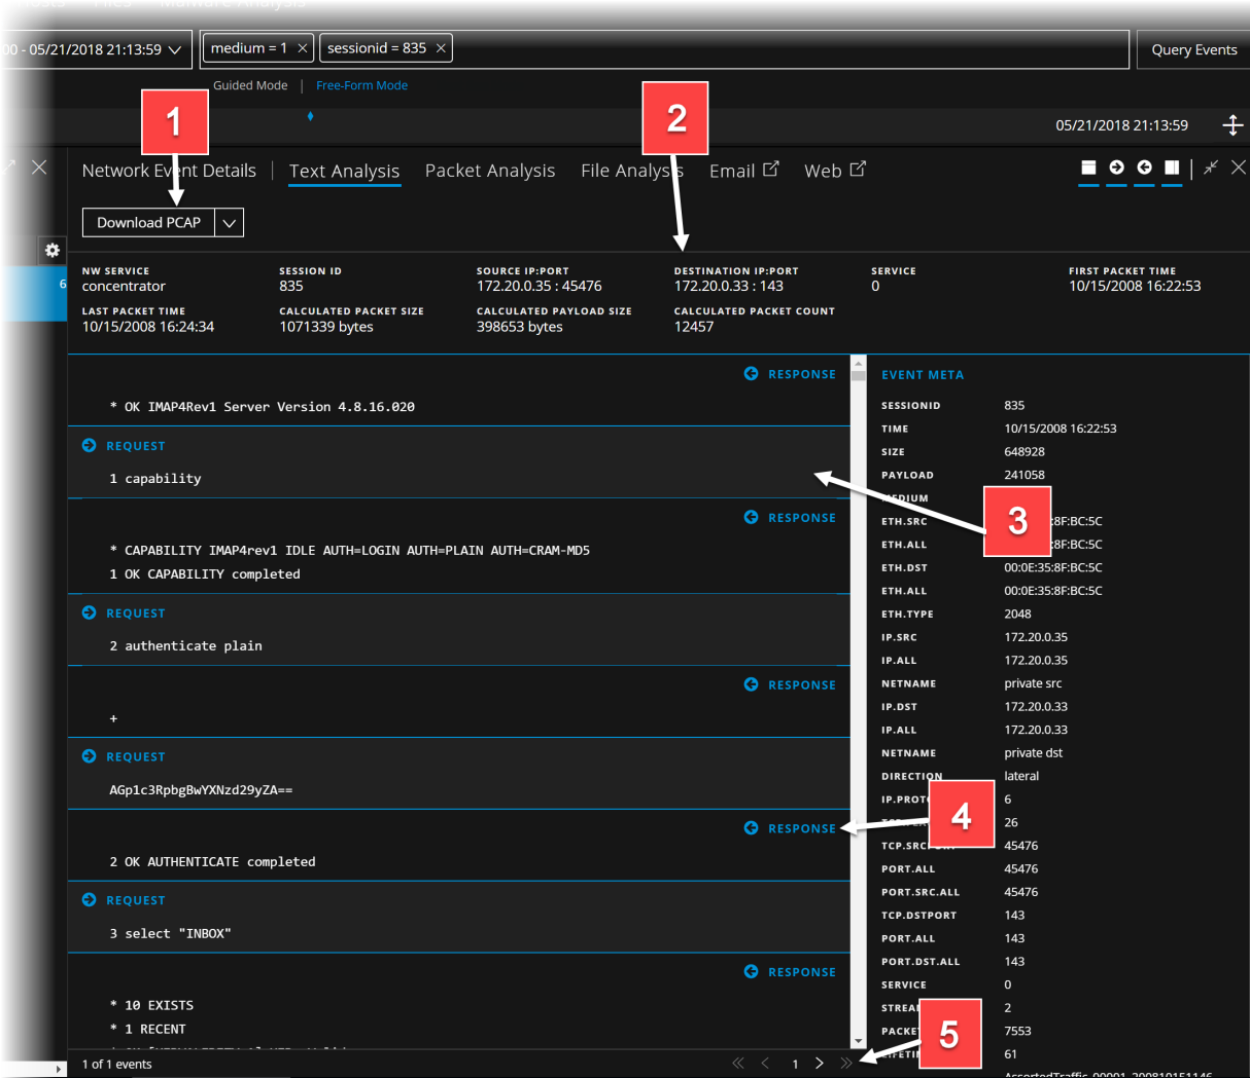
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View - Packet Tab
- Events View - Text Tab
- Events View - File Tab
- Events View - Email Tab
- Events View - Host Tab

Quick Look

The Events view displays the text of a single event in the Text panel (formerly known as Text Analysis). When you click an event in the Event list panel, the adjacent panel shows the text reconstruction. Only the raw log for log events and endpoint events is shown in the Text panel. For network events, the direction of the packet (Request or Response) and contents of each packet are provided in text format. For more examples of the Text, see Reconstructing and Analyzing Events. For detailed procedures, see Analyze Events in the Events View.

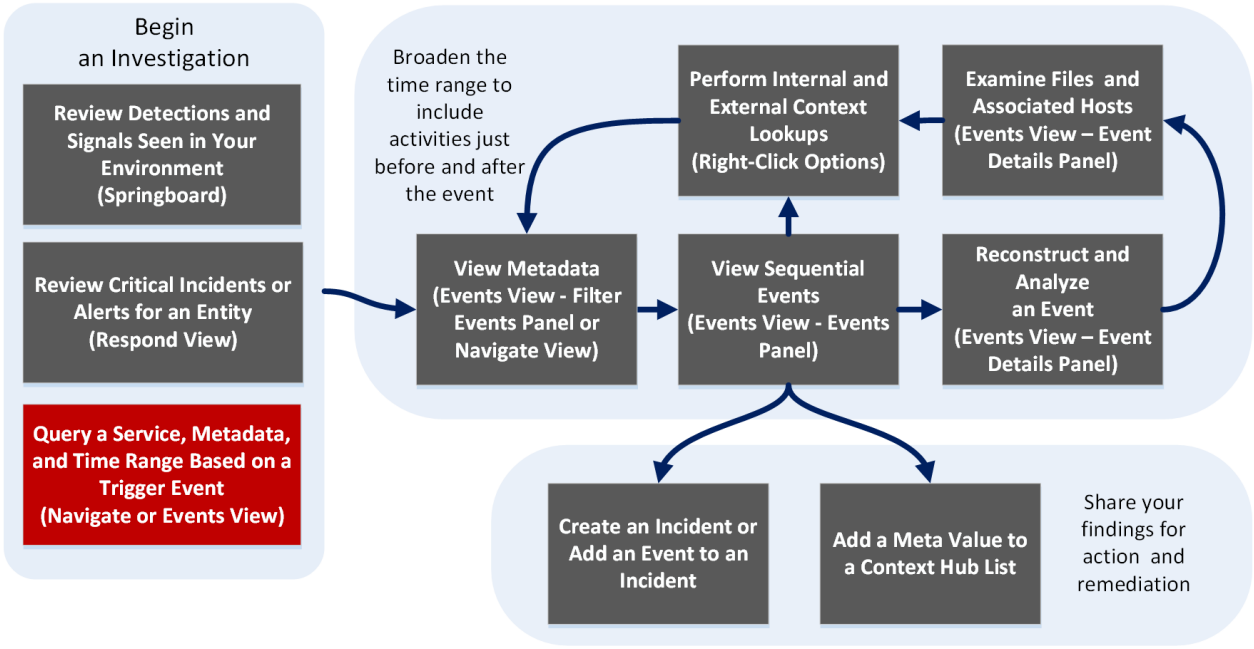


1	Options for exporting a log, a PCAP, or files for deeper analysis and to share with others. This download menu is for network data.
2	The event header information.
3	The payload for a network event includes requests and responses. This is the request side of the packet.
4	This is the response side of the packet.
5	(Version 11.2 and later) Event pagination controls allow more flexibility in paging through a list of events. When a control is unavailable, the image is dimmed; for example, when you are viewing page 1, the ⏪ and ◀ controls are dimmed. ⏪ - Go to the first page ◀ - Go to the previous page ▶ - Go to the next page ⏩ - Go to last page (Only available after last page has already been navigated to)

Investigate Dialog

In the Investigate dialog, analysts can select a service or a collection to investigate. The dialog is automatically displayed when you first go to the Navigate view or Legacy Events view and have not selected a default service to investigate. To access the dialog from a current investigation, select the current service name in the toolbar.

Workflow



What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)

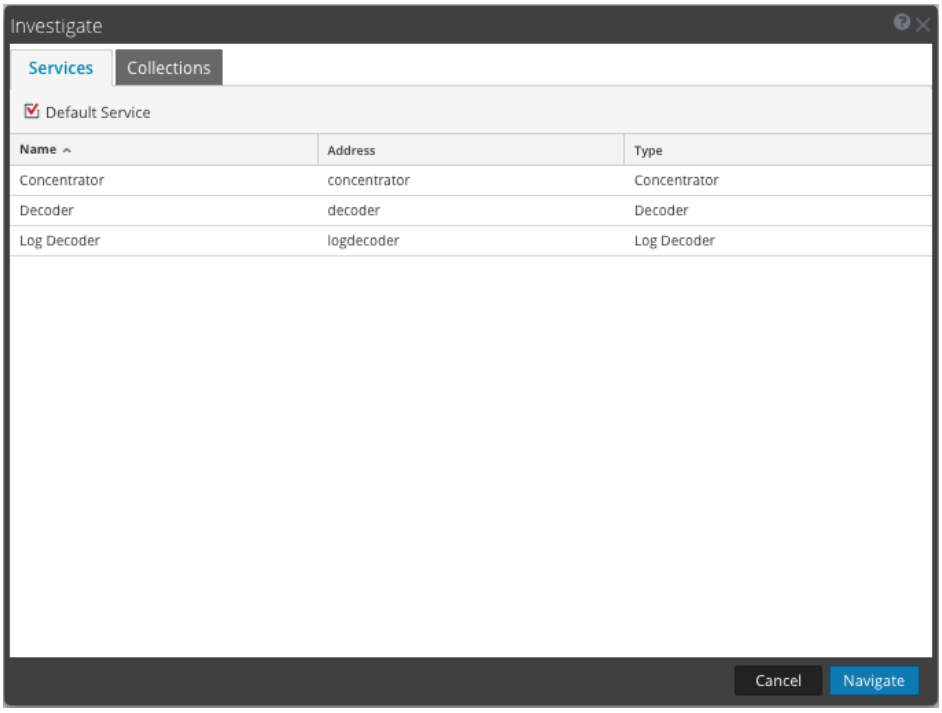
User Role	I want to ...	Show me how
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Navigate View
- Legacy Events View

Quick Look



The Investigate dialog has two tabs: Services and Collections.

Note: Collections are also known as workbench collections. You can only view workbench collections that you have created, and only administrators can create a workbench collection.

The Services tab includes a list of services available for investigation, and three buttons. All features are described in the following table.

Feature	Description
Default Service	Clicking this button sets or clears the default service to investigate. When a service has been set as the default service, the word (Default) is appended to the service name.
Name	The name of the service.
Address	The IP address of the service.
Type	The type of service.
Cancel	Closes the dialog.
Navigate	Opens the selected service in the Navigate or Legacy Events view.

The Collections tab has two buttons and two panels: Workbench and Collections.

The Workbench panel lists available Workbench services by name. After a Workbench service is selected, you can select a collection from the Collections panel.

The Collections panel lists available collections to investigate. After a collection is selected, you can click Navigate to view the collection.

The following table describes the features of the Collections panel.

Feature	Description
Name	The name of the collection.
Type	The type of collection.
Size	The size of the collection.
Data Type	The type of data within the collection.
Date Created	The date the collection was created.

Investigation Tab - User Preferences Panel

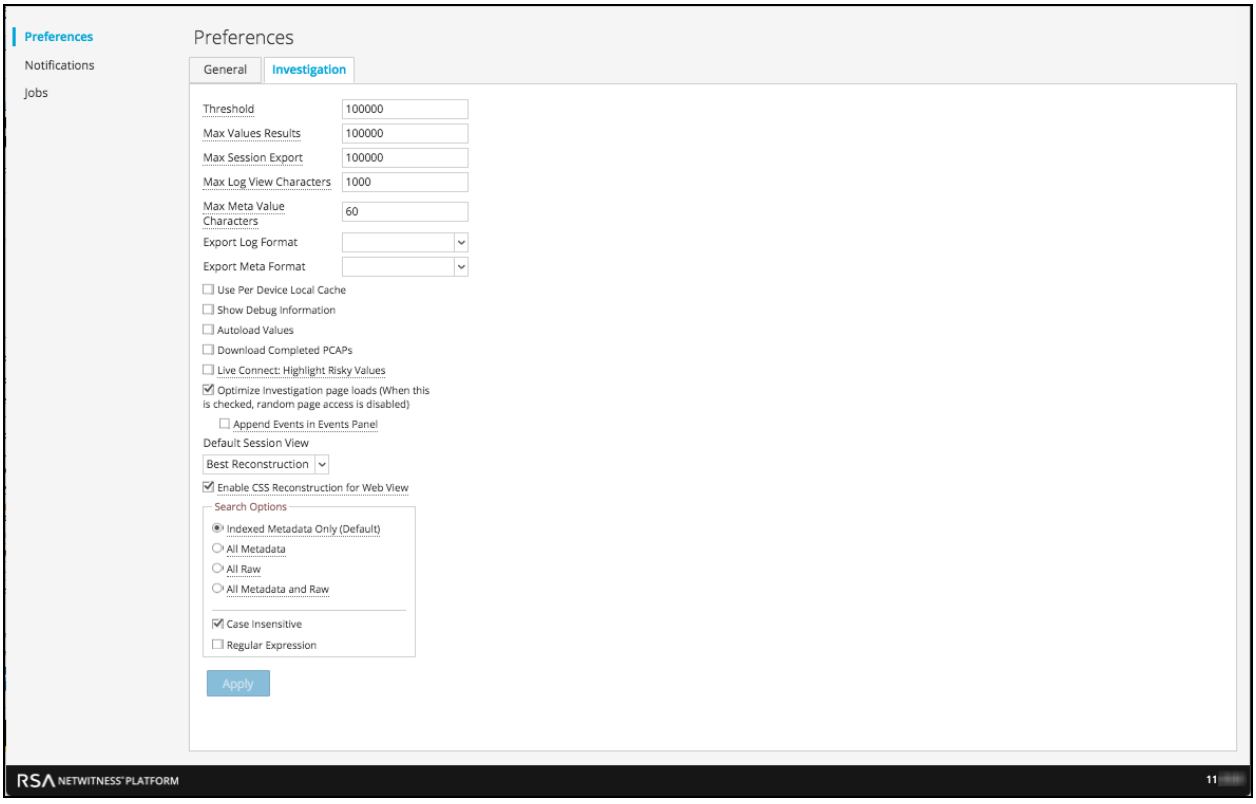
In the Profile view > Preferences panel > Investigation tab, users can set several preferences that affect the performance and behavior of NetWitness when analyzing data, viewing events, and reconstructing events in NetWitness Investigate. To access this tab, select > Profile from the Navigate view or the Legacy Events view. When the Profile view is displayed, select **Preferences > Investigation**. You can change user preferences at any time when you are working in NetWitness.

Related Topics

- How NetWitness Investigate Works
- Navigate View
- Legacy Events View

Quick Look

This figure is an example of the Investigation tab, and the following table describes the preferences that affect Investigate. There are slight differences between the 11.1 search settings and later versions of the search settings and these are explained in Search for Text Patterns in the Navigate and Legacy Events Views.



Feature	Description
Threshold	This setting controls the count shown for a meta key value in the Navigate view during the load. A higher threshold allows more accurate counts for a value. However, a higher threshold causes longer load times. When the threshold is reached, NetWitness displays the count and the percentage of time used to reach the count in comparison to the time necessary to load all sessions with that value. For example, **(>100000 - 18%)** indicates that the threshold was set at 100000 and this load took only 18% of the time it would have taken with no threshold set. The default value is **100000**.
Max Values Results	This setting controls the maximum number of values to load in the Navigate View when the Max Results option is selected in the Meta Key Menu for an open meta key. The default value is **1000**.
Max Session Export	This setting controls the maximum number of sessions that can be exported. The default value is **100000**.
Max Log View Characters	This setting controls the maximum number of characters to be displayed on **Investigate > Legacy Events > Log Text**. The default value is **1000**.
Export Log Format	This setting specifies the default format for exporting logs from Investigate. Available options are **Text**, **XML**, **CSV**, and **JSON**. There is no default value for the log export format. If you do not select a format for logs here, NetWitness displays a selection dialog when you invoke export of logs. When you select one of the options from the Export Log Format drop-down menu and click Apply, the setting goes into effect immediately.
Export Meta Format	This setting specifies the default format for exporting meta values from Investigate. Available options are Text, XML, CSV, and JSON. There is no default value for the meta export format. If you do not select a format for exporting meta values here, NetWitness displays a selection dialog when you invoke export of meta values. When you select one of the options from the Export Meta Format drop-down menu and click Apply, the setting goes into effect immediately. **Note:** If you upgrade to version 11.5.2, the **Export Meta Format** preference is not retained and is reset to blank. You must re-configure this value after you upgrade to version 11.5.2.
Use Per Device Local Cache	Allows you to specify the use of locally cached data from the selected service. By default, this checkbox is cleared (Off), which means that Investigate sends a fresh query to the database rather than displaying cached data in the Investigate views after the initial load. If the option is set, Investigate uses the data from local cache.
Show Debug Information	When this option is set, NetWitness displays the `where` clause beneath the breadcrumb in the Navigate view. For each meta value load, the load time is displayed. If the service is a Broker, the elapsed time for each aggregated service is reported. The default value is **Off.**

Feature	Description
Append Events in Events Panel	When this option is set, the events displayed in the **Events Panel** are added incrementally rather than overwriting the currently displayed events. Each time you click the next page icon, the additional events are appended to the previous events; 1 -25, then 1 -50, then 1 -75 and so on. **Note:** This option is available only if the Optimize Investigation Page Loads option is enabled.
Autoload Values	When this option is set, the service values are automatically loaded in the Navigate view. When not set, NetWitness displays a **Load Values** button, allowing the user the opportunity to modify the options. The default value is **Off**.
Download Completed PCAPs	This setting automates the downloading of extracted PCAPs in the Investigate so that you do not have to manually download and open extracted PCAP files in an application, such as Wireshark, that can handle viewing data in a PCAP format.
Live Connect: Highlight Risky Values	If you want NetWitness Platform to highlight and display only IP addresses that are considered to be risky by the RSA community, set this option. When not enabled, NetWitness Platform displays all IP addresses. By default, this option is cleared (Off).
Optimize Investigation Page Loads	This option is enabled by default (checked) and controls how the Legacy Events view retrieves events. When enabled, results are returned as quickly as possible, but you cannot go to a specific page in the event list. Clearing the checkbox changes the Events list pagination to allow you to go to a specific page in the list (or to the last page). Being able to go to any page in the list costs additional overhead to determine the events in advance.
Default Session View	This setting selects the default reconstruction type for the initial reconstruction view. By default events are reconstructed using the reconstruction type most appropriate to the event.
Enable CSS Reconstruction for Web View	This setting controls how web content reconstruction is performed. If enabled, the web reconstruction includes cascaded style sheet (CSS) styles and images so that its appearance matches the original view in a web browser. This includes scanning and reconstructing related events, and searching for stylesheets and images used in the target event. The option is enabled by default. Clear the checkbox if there are problems viewing specific websites. **Note:** The appearance of the reconstructed content may not match the original web page perfectly if related images and stylesheets could not be found or were loaded from the web browser's cache. Also, any layout or styling that is performed dynamically via client side javascript does not render in the reconstruction because all client side javascript is removed for security purposes.
Search Options	This setting specifies the default search options to apply to a search in the Navigate and Legacy Events views. Search for Text Patterns in the Navigate and Legacy Events Views provides detailed information.
Apply	Saves your preferences and puts them into effect immediately.

Investigate View

The Investigate view is the primary entry point to NetWitness Investigate. In Version 11.5, several of the Investigate submenus are moved to the main menu for easier access. Prior to Version 11.5, the Investigate view had six submenus, which opened different views that allow you to analyze events from different perspectives. The submenus that remain under Investigate are: Navigate, Legacy Events, Events (formerly Event Analysis), and Malware Analysis. The Hosts, Files, and Users views (formerly Entitites) are available from the main menus to improve the workflow for analysts.

Note: The Legacy Events view was the original Events view (11.0 to 11.3.x.x). In Version 11.4 and later, the Legacy Events is no longer needed and it is hidden unless the administrator enables it. By default only the Events view appears in the menu, but when the Legacy Events view is enabled, both the Events view and the Legacy Events view are visible in the menu bar.

How NetWitness Investigate Works provides an introduction to all of the capabilities available in the Investigate view.

Legacy Event Reconstruction View

The Event Reconstruction view is deprecated in favor of the Events view. The Legacy Events view provides a reconstruction of a selected event from the Legacy Events view. By default, NetWitness displays the best reconstruction for the event determined by the event content, or the default reconstruction that you have selected in the Default Session View setting for Investigate. You can use the options in the Event Reconstruction toolbar to change the reconstruction method, view top-to-bottom or side-by-side results, select request and response views, export an event, export meta values, extract files, open an email attachment, and open the event in a new tab.

To access this view, do one of the following:

- In any Legacy Events view, double-click an event.
- In the Legacy Events view with Detail View selected, right-click **Events** at the end of the event, and select **Event Reconstruction.**
- In the Event Reconstruction toolbar of previewed reconstruction, click **Open Event in New Tab**.
- In the Navigate view, select **Actions > Go to event in Event Reconstruction**, and enter an event ID.

What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)
Threat Hunter	view sequential events*	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View

User Role	I want to ...	Show me how
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident*	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

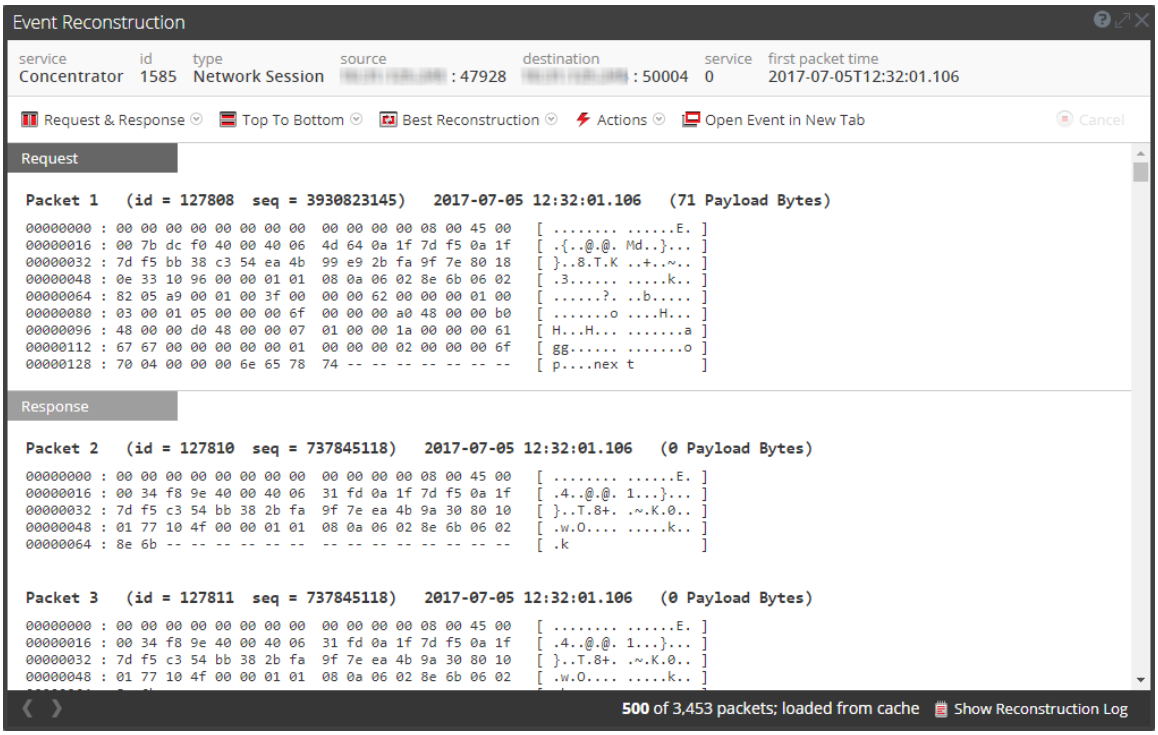
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Events View

Quick Look

This figure is an example of the Event Reconstruction view. The following table describes the toolbar options.



Feature	Description
Request & Response	Displays a drop-down menu for selecting whether the view displays: • Request & Response • Request • Response
Organization	Displays a drop-down menu for selecting whether the information is displayed top to bottom or side by side.
Reconstruction View	Displays a drop-down menu for selecting what information is displayed. By default, Best Reconstruction is selected. Other options are: • View Meta • View Text • View Hex • View Packets • View Web • View Mail • View Files
Actions	Displays a drop-down menu with the actions available in the Event Reconstruction view (Export PCAP, Extract Files, and Export Meta).

Feature	Description
Open Event in New Tab	Opens the event in a new browser tab.
Event Analysis	Open the event in the Event Analysis view.

Beneath the toolbar is a list of meta keys and values. Some of the keys offer a drop-down menu with available actions.

The bar at the bottom of the view offers several options.

Feature	Description
	Displays the previous event.
	Displays the next event.
Show Reconstruction Log	Displays the reconstruction log at the bottom of the view. Once you click this button, it changes to Hide Reconstruction Log.

Legacy Events View

The Legacy Events view is deprecated in favor of the Events view. In the Legacy Events view a list of events associated with a session is available; this view is optimized for viewing raw events in sequence by time. You can display the events list in several forms, filter events, search for events, and open a reconstruction of an event.

There are two ways to display the Legacy Events view:

- Go to **Investigate > Legacy Events.** NetWitness runs a default query on the last three hours for the default service (if one is set) or displays a dialog in which you can select a service and then runs the default query. The default query selects all events and the Legacy Events view displays events on the selected service, with the oldest events first.
- From within the **Navigate** view, double-click an event. The Legacy Events view displays the events on the selected service based on the drill point in the Navigate view.

Note: The Legacy Events view was the original Events view (11.0 to 11.3.x.x). The Legacy Events is no longer needed and it is hidden unless the administrator enables it. By default only the Events view appears in the menu, but when the Legacy Events view is enabled, both the Events view and the Legacy Events view are visible in the menu bar.

What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)
Threat Hunter	view sequential events*	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event*	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View

User Role	I want to ...	Show me how
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident*	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Filter Results in the Legacy Events View
- Downloading and Acting Upon Results

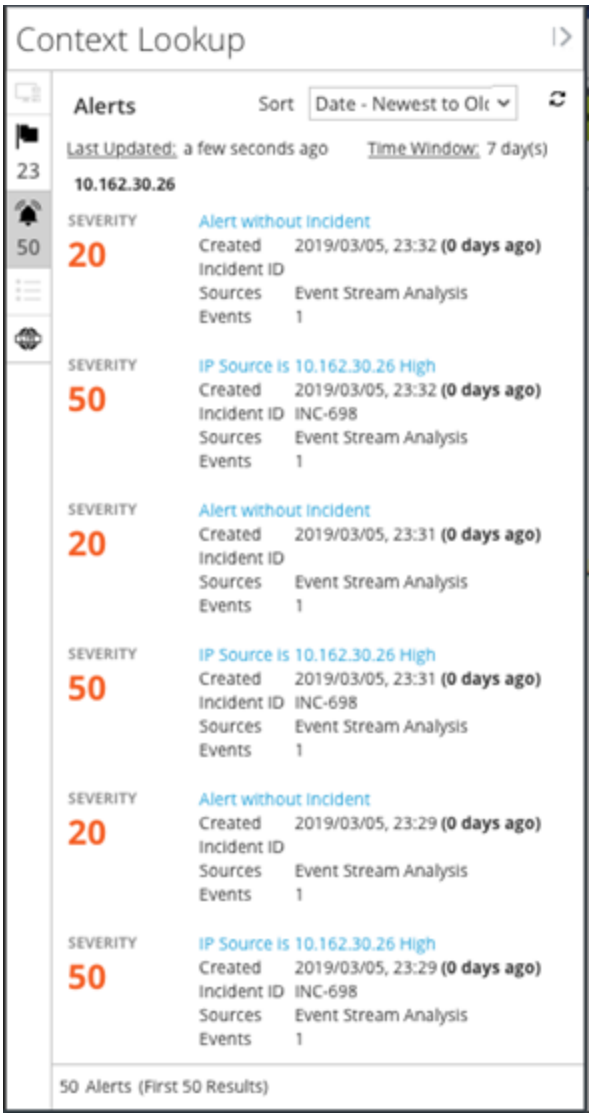
Quick Look

The Legacy Events view provides three built-in presentations of event data: the Detail view, the List view, and the Log view. The List view and Detail view provide more information for each event including the timestamp, event type, event theme, and size.

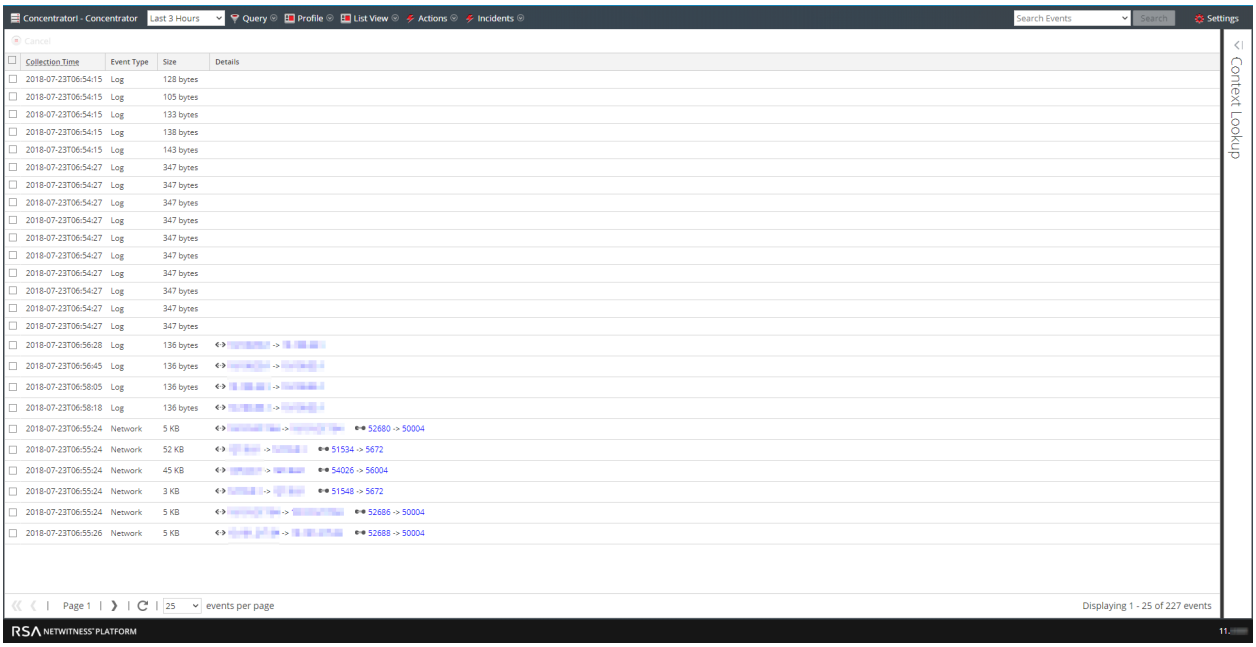
- The List View shows corresponding source and destination address and port information for events in summary form in a grid.
- The Detail View shows all metadata collected for the event in a paged view.
- The Log View is optimized for viewing log and endpoint information, and provides more information for each log including the timestamp, event type, service type, service class, and the logs.

You can use queries, the time range setting, and profiles to filter the events listed in the Legacy Events view. From any view type in Legacy Events view, you can extract files; export network events, endpoint events, logs, and meta values, and open the Event Reconstruction panel. In the Detail View you can also open the event in the Events view.

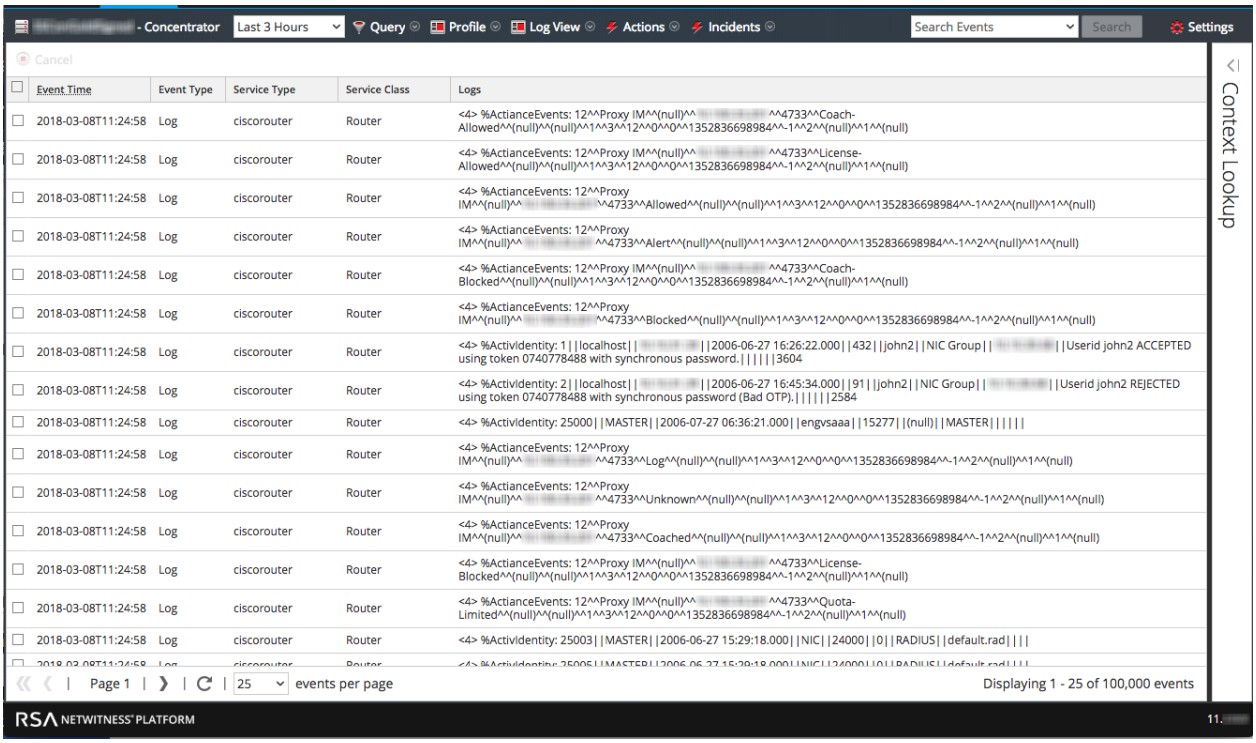
The following figure is an example of events in the Detail View. The Context Lookup panel is visible only if the Context Hub service is configured.



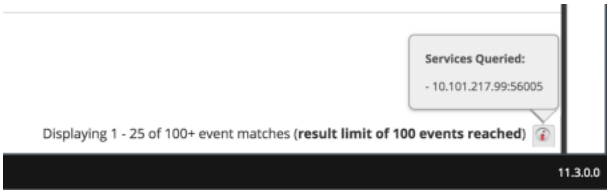
The following figure is an example of events in the List View.



The following figure is an example of the Log View.



The following figure shows the information added to the footer for Version 11.3 and later.



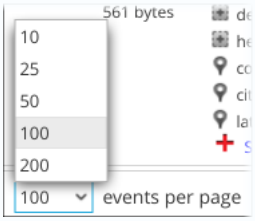
Detailed Description

The Legacy Events view has a toolbar at the top with the following options.

Feature	Description
Select Service	Displays the selected service name next to the icon. Opens the Investigate dialog, in which you can select a service for which the event list is displayed.
Time Range	Displays a drop-down menu for selecting the time range to apply to the event list. You can choose one of the standard options or specify a custom time range.
Query	Displays the Query dialog, in which you can enter a custom query directly instead of drilling down the data (see Create a Query in the Navigate and Legacy Events Views).
Profile	Displays the Profile menu; the currently selected profile is displayed in the toolbar. The menu options include built-in (Default) profiles and custom profiles, as well as an option to manage profiles. Each profile can include a meta group, a column group, and a beginning query that is applied to the Navigate view (meta groups and queries) and the Legacy Events view (column groups and queries) as you investigate events. (see Use Query Profiles to Encapsulate Common Areas for Investigation).
View Type Drop-down	Displays a drop-down menu for selecting the event view type. • Detail View shows events in a paged format with detailed information for each event. • List view shows the events in table form with a summary of each event in a separate row. • Log View shows a log-oriented events grid with a summary of each log in a separate row. • Custom Column Groups displays the event list using a column group selected from a drop-down list of custom column groups. • Manage Column Groups displays the dialog for creating and editing custom column groups.
Actions	Displays a drop-down menu with actions in the Legacy Events view: • Export an events as PCAP files, export logs, export endpoint events, or export meta values. • View an event reconstruction in a popup window or in a new tab. • Reset all filters in the Legacy Events view.
Incidents	Create a new incident in Respond and add the selected events, or add selected events to an existing incident in Respond.

Feature	Description
Search	Displays the Search Events options, which allow you to specify the export log and export meta value format with additional options explained in Search for Text Patterns in the Navigate and Legacy Events Views.
Settings	Displays the Investigation settings for the Legacy Events view (which are also available in the Profile view) so that you can change Investigation settings without navigating away from the Legacy Events view. When you change a setting In the Legacy Events view, the setting is also changed in the Profile view (see Configure the Navigate View and Legacy Events View).

Other features of the Legacy Events view are described in this table.

Feature	Description
Show Additional Meta (in the Detail View of an event)	Displays the rest of the metadata for the event.
Event Analysis (in the Detail View of an event)	Opens the selected event in the Events view.
 (in the footer)	Pagination controls allow more flexibility in paging through a list of events. When a control is unavailable, the image is dimmed; for example, when you are viewing page 1, the « and ‹ controls are dimmed. « - Go to the first page ‹ - Go to the previous page 3 Page 3 - Go to a specific page › - Go to the next page » - Go to the last page  - Select the number of packets per page When you select a number of events per page, the setting is saved in browser cache so that you do not have to select your preferred number of events each time you log in. The setting applies to all views: Log View, List View, and Details View.

Feature	Description
Displaying 1-100 of 100,000 events (in the footer) Displaying 1-25 of 100+ event matches (result limit of 100 events reached) (in the footer)	Displays the count of events displayed versus the total number of matching events. In Version 11.3 and later, the footer includes a notification if the results limit configured by the administrator has been reached to let you know that more results are available but not viewable. To view the additional results, you need to refine the filter to get fewer results. Clicking the information icon  in the footer displays the IP address and connecting port number for all services queried.

Manage Default Meta Keys Dialog

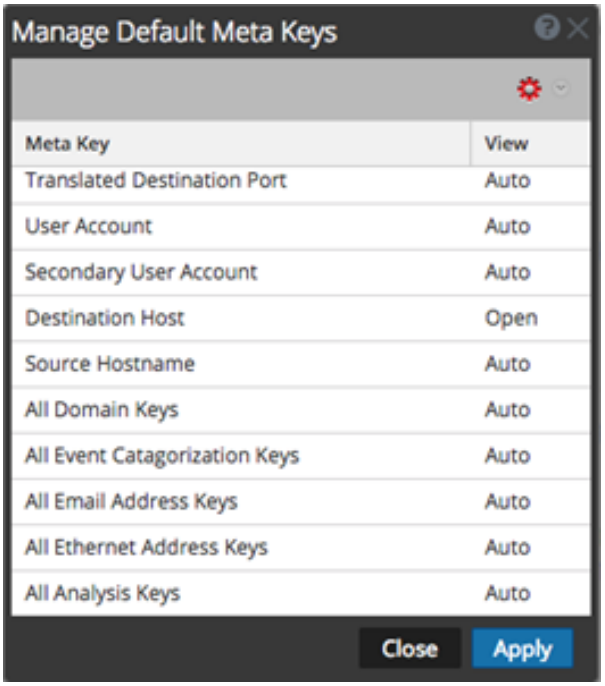
In the Manage Default Meta Keys dialog, analysts can specify the meta keys to be displayed in the Navigate view > Values panel (see Manage and Apply Default Meta Keys in an Investigation. This can help you find the desired data more quickly and prevents the loading of meta keys that are not of interest. To access this dialog, in the **Navigate View** toolbar, select **Meta > Manage Default Meta Keys**.

Related Topics

- How NetWitness Investigate Works
- Use Meta Groups to Focus on Relevant Meta Keys

Quick Look

The following figure illustrates the Manage Default Meta Keys dialog, which has a list of meta keys, toolbar, Close button, and Apply button. In the list, you can view, sort, and manage default meta keys. If you click and drag meta keys, you can rearrange their order. The following table describes columns in the list.



Column	Description
Meta Key	This column displays the meta keys available for the service. In Version 11.1 and later, default meta entities are also included, for example All Domain Keys and All Email Address Keys.

Column	Description
View	This column displays the type of view assigned to each meta key. By clicking on the view in each row, you can assign the meta key a different default view. There are four views: • Auto: Reverts to the default view for meta keys as specified in the service index file. • Close: The values of this meta key are closed by default, and can be opened manually. • Hidden: These meta keys are hidden by default, and are not shown in Investigation at all. • Open: The values of this meta key are displayed by default. When you modify the default meta keys for a non-indexed meta key, you cannot set the key to **Open**. If you change the default view for a group of meta keys to **Open** and some of the meta keys are non-indexed, the non-indexed meta keys revert to **Auto**. As a result, the meta key is automatically loaded only if it is indexed, and non-indexed meta keys are **Closed** until opened manually.

The following table describes the toolbar options and buttons.

Feature	Description
	Displays a drop-down menu that allows you change the default view of all the meta keys. There are four views: • Auto: Reverts to the default view for meta keys as specified in the service index file. • Close: The values of this meta key are closed by default. • Hidden: The values of this meta key are hidden by default. • Open: The values of this meta key are displayed by default.
Close	Closes the dialog. Any unsaved changes are lost.
Apply	Applies the changes, and they become effective immediately.

Meta Groups Dialogs

You can use meta groups to filter data displayed in an investigation. A fresh installation of NetWitness includes built-in meta groups to help you find interesting data sets in Investigate. The built-in meta groups are prefixed with RSA for identification and can be duplicated but cannot be edited or deleted. You can create your own groups and you can duplicate and edit a built-in group to create a custom group. With a meta group in effect during an investigation, the information in the Navigate view and Events view includes only the meta keys in the selected group.

While the functionality of meta groups is similar in the Navigate view and the Events view, the user interface and some of the procedures are different.

Using options in the Events view Meta Groups menu (Version 11.5 and later) , you can:

- Select a meta group to apply.
- See the details of a meta group.
- Create, edit, and delete custom meta groups.
- Clone and edit the clone of a built-in or custom meta group.

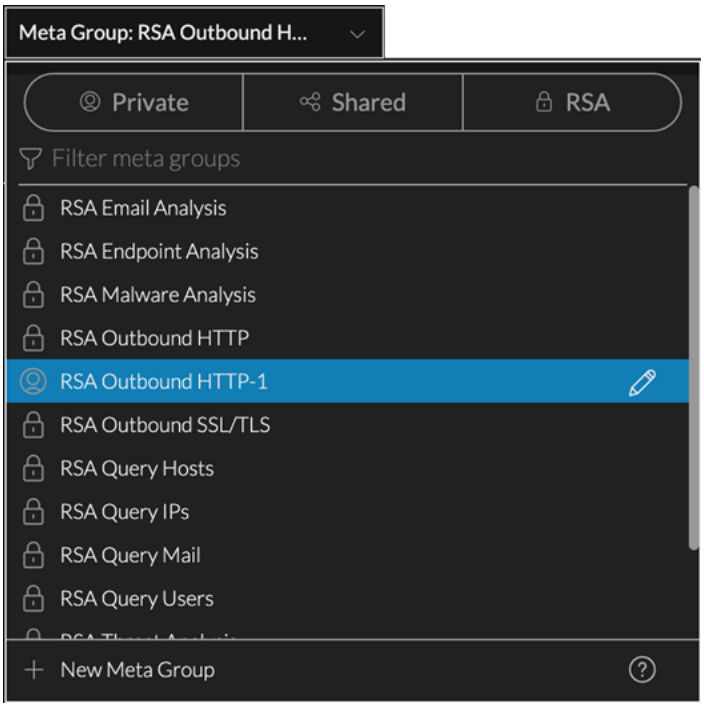
Using options In the Navigate view Manage Meta Groups dialog, you can do all of the above as well as import and export a meta group. Refer to Use Meta Groups to Focus on Relevant Meta Keys for detailed information.

Related Topics

- How NetWitness Investigate Works
- Use Meta Groups to Focus on Relevant Meta Keys
- Filter Results in the Navigate View

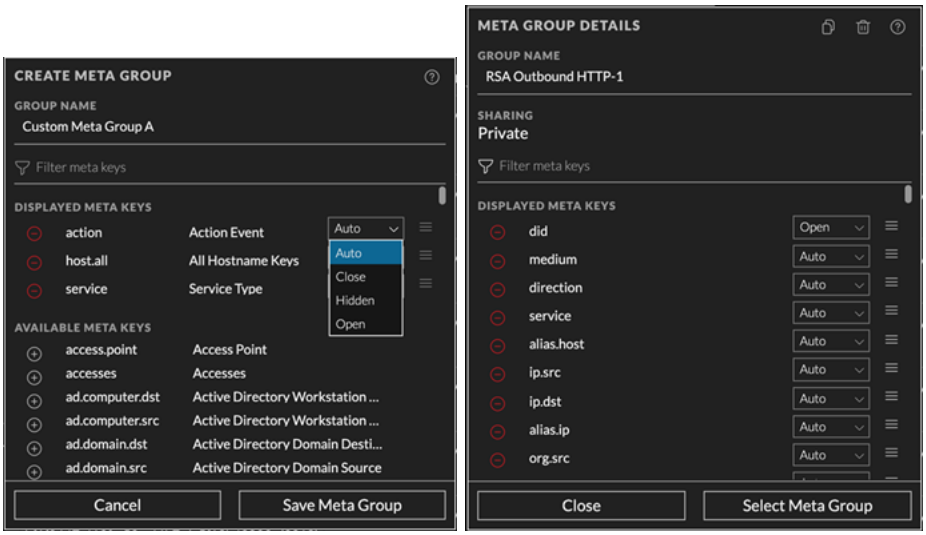
Quick Look - Meta Groups Menu, Create Meta Group Dialog, and Meta Group Details Dialog

This section introduces the Meta Groups Menu, Create Meta Group dialog, and the Meta Group Details dialog. The following figure is an example of the Meta Group menu. The table describes the options.



Feature	Description
Visibility Options	Control the types of meta groups that are visible in the list. You can use any combination of the visibility options: **Private**, **Shared**, or **RSA** (blue = selected, black = not selected). Initially none of the buttons are selected and all meta group types are visible. This is the same result as if all three buttons are selected. The visibility options work together with text in the Filter Meta Groups field. If the visibility option is hiding built-in groups (which include "RSA" in the group name) and you search for a name that contains "RSA," the list is empty. Private = display private groups that only you can manage Shared = display shared groups that anyone in your organization can manage RSA = display built-in groups that only RSA can manage
Filter Meta Groups	Filters the list of meta groups as you type text so that only group names that contain the typed text are displayed.
Meta Group List	The list of meta groups consists of custom and built-in groups. Custom meta groups can be shared or private. The RSA meta groups are built-in meta groups; you cannot edit or delete these, but you can make a copy and edit the copy. Icons preceding the meta group name distinguish the private groups, shared groups, and built-in groups.
New Meta Group	Displays the Create Meta Group dialog, where you can create a custom meta group.

The Create Meta Group dialog, shown in the figure on the left, allows you to define a custom meta group. The figure on the right illustrates the Meta Group Details dialog, in which you can edit a custom meta group. The table describes the fields and options in the dialogs.

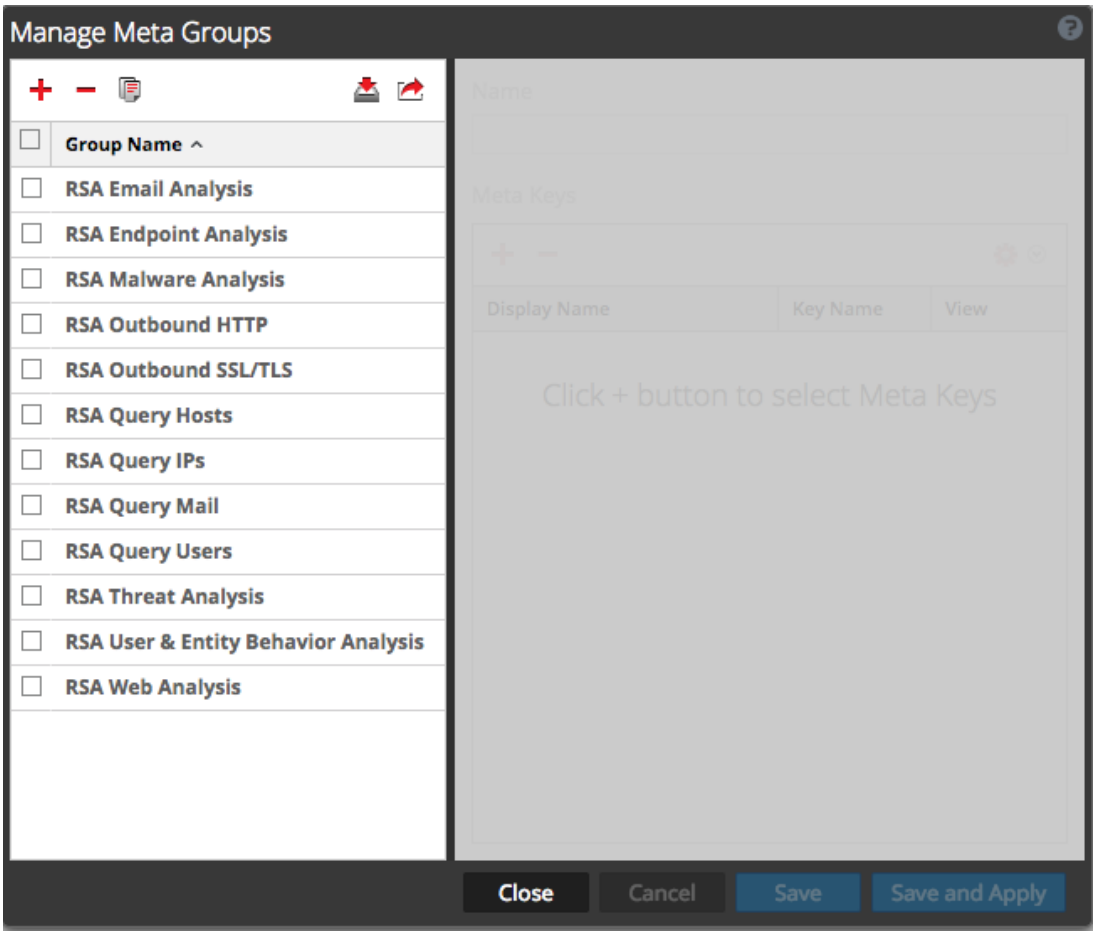


Feature	Description
	Creates a copy of the meta group so that you can edit a copy. This is useful if you want your own copy of a built-in group, a shared copy of a private group, or a private copy of a shared group.
	Deletes the custom meta group that you are currently editing. This action is irreversible and applies globally. If the meta group is a shared group, it is no longer available to anyone.
Group Name	Displays the name of the meta group. The name must be unique and contain fewer than 64 characters. You can type in this field to edit the name in a custom meta group.
Sharing	Specifies whether the meta group is shared or private. This setting is available when you first create the group. After it is created, you cannot change a shared column group to private, or a private column group to shared.
Filter Meta Keys	Filters the Displayed Meta Keys and Available Meta Keys based on the text that you type. Only meta keys that contain the typed text are displayed.
Displayed Meta Keys	Displays a scrollable list of meta keys that are selected for use in the custom meta group. You can add meta keys in the Available Meta Keys list to this list, remove meta keys from this list , and drag meta keys up or down to change the order in this list ( ). Drag and Drop is disabled when text is typed in the Filter Meta Keys field. For each displayed meta key you can choose:
Available Meta Keys	Displays a scrollable list of meta keys that are available (on the service) for use in the custom column group. You can add them to the Available Meta Keys list. Clicking next to the meta key name adds it to the Displayed Meta Keys list. You can also set the initial view of each meta key: Open, Closed, Hidden, or Auto (the default setting).

Feature	Description
Initial View Option	For each meta key, you can set the initial view option: -When set to Auto, the meta key is automatically loaded only if it is indexed, and non-indexed meta keys are Closed until opened manually. If you change the default view for a group of meta keys to Open and some of the meta keys are non-indexed, the non-indexed meta keys revert to Auto. -Open meta keys are listed in the Filter Events panel, and the values are loaded. -Closed meta keys are listed in the Filter Events panel, but the meta values are not loaded until you open the meta key. -Hidden meta keys are not listed in the Filter Events panel at all. This is useful if you are using a single meta group for multiple purposes instead of creating several meta groups; you can turn off certain keys off without removing them from the meta group. You can also use the Hidden view when testing out some new keys or if you want to prepare a meta group with some new meta keys that are not yet available and would error out if in an Auto, Open, or Closed state.
	Allows you to drag and drop meta keys in the Displayed Meta Keys list so that you can see the data in the order you prefer.
Close button	Closes the dialog.
Save Meta Group	For the Create Meta Group dialog only, saves the new meta group.
Reset	For the Meta Group Details dialog only reverts the edited meta group to the last saved state.
Update Meta Group	For the Meta Group Details dialog only, applies changes to an edited meta group.
Select Meta Group	Applies the meta group. The Filter Events panel is refreshed to display only the meta keys in the selected meta group.

Quick Look - Manage Meta Groups Dialog

The following figure is an example of the Manage Meta Groups dialog.



The Meta Groups panel is on the left side of the Manage Meta Groups dialog. This is where you can add, delete, import, and export meta groups. The following table describes the features of the Meta Groups panel.

Feature	Description
	Adds a meta group using the Settings panel on the right side of the Manage Meta Groups dialog.
	Deletes the selected meta group. A confirmation dialog is displayed before the meta group is deleted.
	Creates a copy of the selected meta group.
	Displays the Meta Group Import dialog, where you can upload a file.
	Exports the selected meta group to your computer.
Group Name	Lists all meta group names.

The Settings panel is on the right side of the Manage Meta Groups dialog. This is where you create and edit meta groups. Below the Name field is the Meta Keys list. The following table describes the features of the Settings panel.

Feature	Description
Name	Displays the name of the selected meta group.
	Displays the Available Meta Keys dialog, where you can select meta keys to add to the group.
	Deletes the selected meta keys.
	Displays a drop-down menu, where you can select the view for all meta keys. There are four options based on the possible values for the `defaultAction` property used to define a key in the custom index file for the service: • Hidden: These meta keys are hidden by default, and are not shown in Investigation at all. • Open: The values of this meta key are displayed by default. • Close: The values of this meta key are closed by default, and can be opened manually. • Auto: Reverts to the default view for meta keys as specified in the service index file.
Display Name	Indicates the name that is displayed for the key in Investigate views, and is defined by the `description` property for the key in the custom index file for the service..
Key Name	Indicates the `name` of the meta key as defined in the custom index file for the service.
View	Indicates the view to which the meta key is set. You can change: • Change the view for all meta keys by clicking  in the View column header, then selecting a view from the drop-down menu. • Change the view for a single meta key by clicking a single meta key in the View column, then  and selecting a view from the drop-down menu.

The following table describes the buttons at the bottom of the dialog.

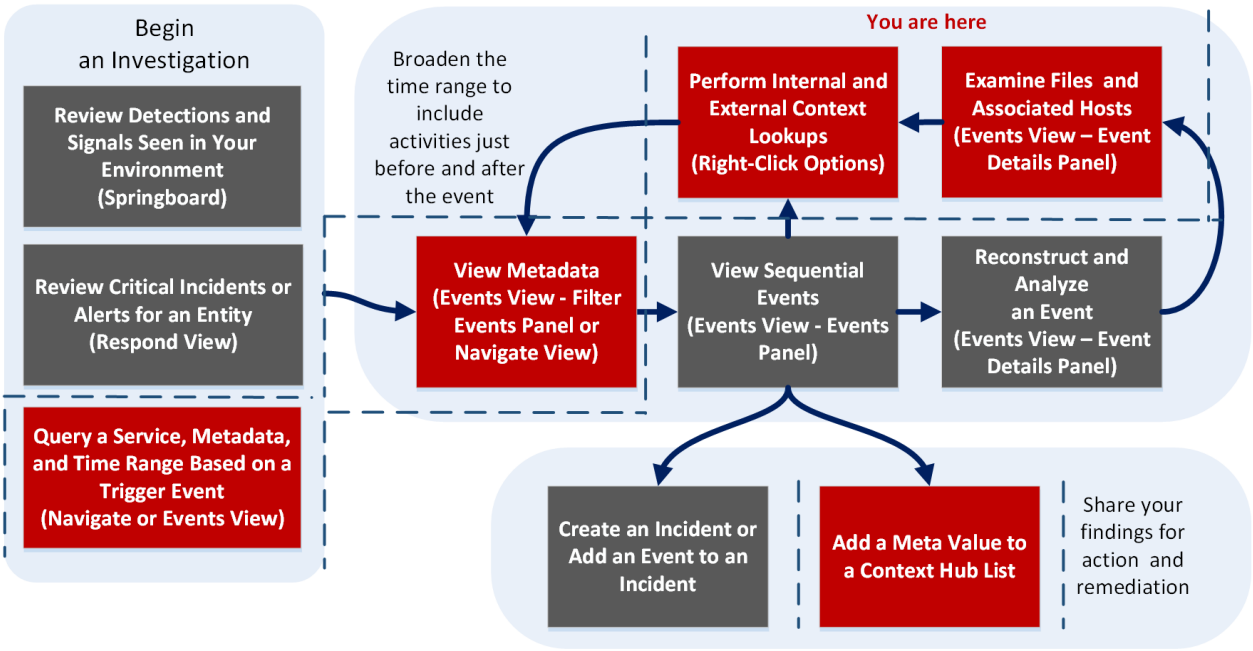
Feature	Description
Close	Closes the dialog.
Cancel	Cancels all changes.
Save	Saves all changes.
Save and Apply	Saves and immediately applies all changes.

Navigate View

The Navigate view (**Investigate** > Navigate) displays event metadata--the meta keys and meta values--that were found in captured data for the selected service. The data is filtered and displayed in accordance with the options you set for profile, time range, meta group, and query. You can also drill into the data by clicking meta keys and meta values.

Note: By default, the Navigate view is disabled in Version 11.6 as the Filter Events Panel in the Events view provides this functionality. To enable the Navigate view, see Configure the Navigate View and Legacy Events View.

Workflow



In the Navigate view, you can:

- View metadata for events in the Values panel.
- Visualize events in a timeline or parallel coordinates chart.
- Save events, go to an event using the event ID, visualize an event, and print an event.
- View additional contextual data for meta keys and values.
- Open a drill point or an event in the Legacy Events or the Events view.

What do you want to do?

User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata*	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts*	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups*	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list*	Look Up Additional Context for Results

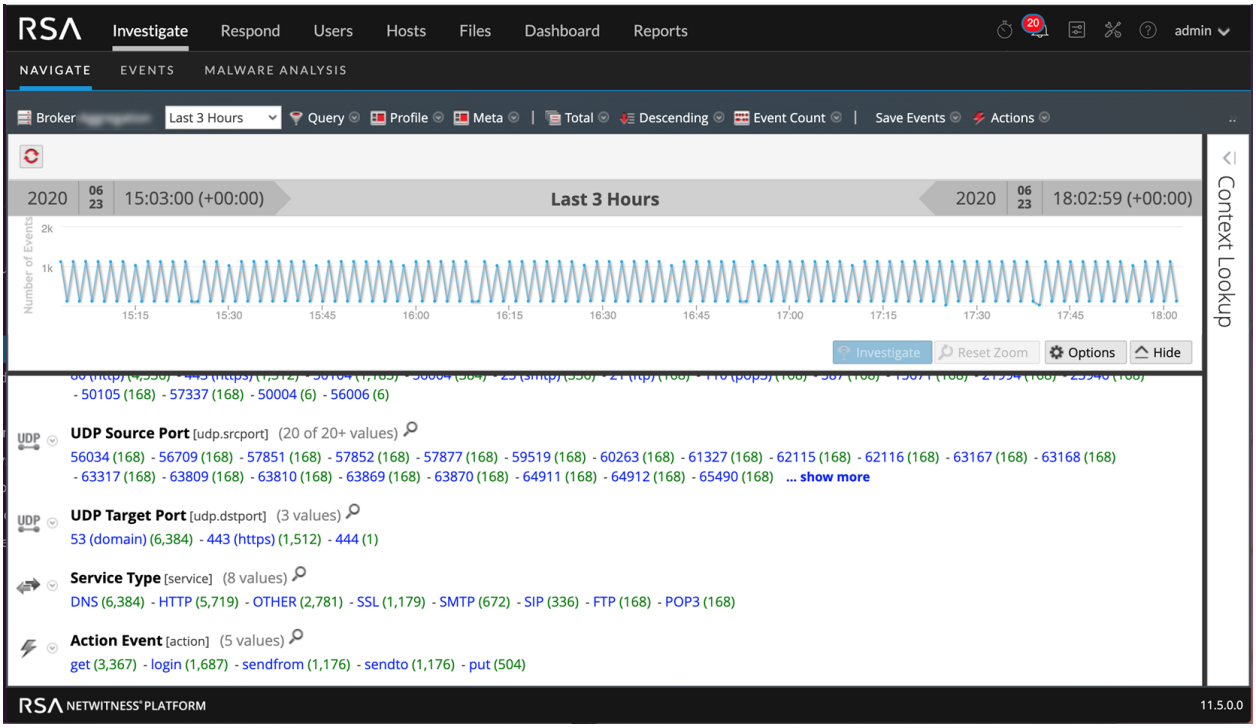
*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Legacy Events View
- Events View

Quick Look

This figure illustrates the Version 11.5 Navigate view.



The Navigate view consists of these features:

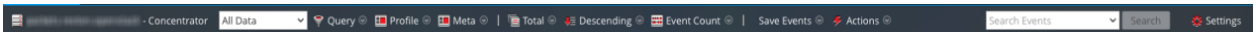
- Toolbar
- Pause/reload button and breadcrumb
- Time banner
- Optional debug information.
- Collapsible Visualization panel
- Values panel
- Context Lookup panel
- Context menus

Toolbar

The following figure is an example of the toolbar. The toolbar provides a way to:

- Change the service being investigated.
- Control the range of data displayed: You can select use profiles, set a time range, use meta groups, and create queries to apply to the data.

- Set the quantification method and sorting method for data in the Values panel.
- Perform actions on the results. You can export and print results, open an event for which you have an event ID in the Legacy Events view or Events view, and pass a query to Informer.
- Configure Investigate settings without navigating away from the Investigate views.



Some of the toolbar options are labeled with the default value or the selected value rather than displaying the name of the option. For example, the time range option in the example above is labeled **Last 5 Minutes** to reflect the currently selected value. These are the toolbar options.

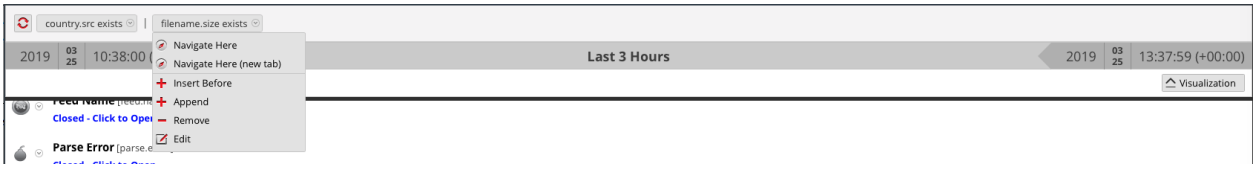
Option	Description
	Displays the selected service name next to the icon. Clicking the icon opens the Investigate a Service dialog, in which you can select a service to investigate and set the default service to investigate (see Begin an Investigation in the Navigate or Legacy Events View). Changing the service does not cause a reload of the data.
Time Range	Displays the Time Range options; the currently selected option is displayed in the toolbar (see Filter Results in the Navigate View). Possible choices are: • All Data • Last 5, 10, 15, or 30 Minutes • Last Hour, Last 3, 6, 12, or 24 Hours • Last 2 or 5 Days • Early Morning • Morning • Afternoon • Evening • All Day • Yesterday • This Week • Last Week • Custom **Note:** If you specify custom start or end times in seconds, the value for start time in seconds always defaults to :00, and the value for end time in seconds always defaults to :59. For example, if you are using time to drill down into an issue, the drill time will be interpreted as `HH:MM:00 - HH:MM:59`. Seconds display in this format in Investigate functions.

Option	Description
Query	Displays the Query dialog, in which you can enter a custom query directly instead of drilling down the data. See Query Dialog for a description of the dialog.
Profile	Displays the Profile menu; the currently selected profile is displayed in the toolbar. A profile allows you to manage and use profiles that can include custom meta groups, a default column group, and a beginning query. The Profiles apply to the Navigate view (meta groups and queries), the Legacy Events view, and the Events view (column groups and queries). See Use Query Profiles to Encapsulate Common Areas for Investigation for more information.
Meta	Displays the Meta Group menu. You can use Default Meta Keys or a custom Meta Group. You also have the option to make changes to both group types (see Use Meta Groups to Focus on Relevant Meta Keys).
Sort Field	Displays the Sort Field menu; the currently selected option is displayed in the toolbar. The menu has two options: Order by Total and Order by Value. The Sort Field is a complement to the Sort Order option; the data for each meta key is ordered based on the total (green number) or the meta value (blue text) (see Filter Results in the Navigate View).
Sort Order	Displays the Sort Order menu; the currently selected option is displayed in the toolbar. The menu has two options: Sort in Ascending Order and Sort in Descending. The Sort Order is a complement to the Sort Field option; the selected sort field for each meta key is ordered in ascending or descending order (see Filter Results in the Navigate View).
Quantification Method	Displays the Quantification Method menu; the currently selected option is displayed in the toolbar. The Quantification Method only applies to the meta key results in the Values panel. It does not apply to the timeline. The drop-down menu contains three options for calculating the quantity (green number in parentheses) for a meta value: Quantify by Event Count, Quantify by Event Size, and Quantify by Packet Count (see Filter Results in the Navigate View). These are applied differently depending on the type of data in view. For packet data: • Quantify by Event Count shows the number of sessions. • Quantify by Event Size shows the size in bytes. • Quantify by Packet Count shows the number of packets. For log data: • Quantify by Event Count shows the number of logs. • Quantify by Event Size shows the size in bytes. • Quantify by Packet Count shows the number of logs.
Save Events	Displays the Save Events menu, in which you can use options to: extract files associated with an event, export the current drill point as a PCAP file, and export the current drill point as a log file (see Export a Drill Point).

Option	Description
Actions	The Actions menu includes actions that you can perform in the Navigate view (see Refining the Results Set). In Version 11.1 and later, the options are Visualize, Go to event in Event Reconstruction, Go to event in Events view, and Print).
Search Events	Enables you to search for text patterns within the current set of events. If you click in the Search field, it shows a drop-down menu with search options. If you click Apply, it saves the selected options and also updates the search options in the Legacy Events view and the Investigations profile (see Search for Text Patterns in the Navigate and Legacy Events Views).
Settings	Displays the settings for the Navigate view (which are also editable in the Profile view) so that you can change Investigate settings without navigating away from the Navigate view. When you change a setting In the Navigate view the setting is also changed in the Profile view (see Configure the Navigate View and Legacy Events View).

Pause/Reload Button and Breadcrumb

The breadcrumb tracks each query as you drill down through the metadata for the service. The following figure is an example of the breadcrumb.



Each query is listed with a drop-down menu in a pipe separated string. The last point is the current point, also called the tip. The icon in front of the breadcrumb allows you to pause the loading of meta values and to reload meta values. The breadcrumb does not include the service name and appears only if a query is in effect. If too many drill points exist for display, the overflow is shown as double angle brackets, >>, at the end of the breadcrumb. Each drop-down menu in the breadcrumb is the same, with slight variation based on the position of the crumb.

The following table describes the controls and menu options in the breadcrumb.

Feature	Description
Pause	Pause and Reload button. Controls the loading of data in the view. It has three possible functions: pause loading, continue loading, and reload.
Navigate Here	Opens the selected drill point in the current Values panel.
Navigate Here (new tab)	Opens the selected drill point in a new tab.
Insert Before	Inserts a query before the current drill point. The Create Filter dialog opens and you can define a custom query to insert in the breadcrumb (see Create a Query in the Navigate and Legacy Events Views).

Feature	Description
Append	Appends a query after the current drill point. The Create Filter dialog opens and you can define a custom query to append to the end of the breadcrumb (see Create a Query in the Navigate and Legacy Events Views).
Remove	Removes the selected drill point from the breadcrumb.
Edit	Opens the selected drill point in the Create Filter dialog so that you can edit the query.
>>	Clicking the angle brackets displays a drop-down menu of the breadcrumb overflow.

(Optional) Debug Information

If you have activated the Show Debug Information setting and the service you are navigating is a Broker, NetWitness, displays the debug information beneath the breadcrumb.

The debug information is the `where` clause from the current query. The only time there is no `where` clause is when the time range is all data and there are no drill points. If the Broker has at least one aggregate service that is offline, the debug information also lists the offline service.

For example:

```
(attachment exists)&&(tcp.dstport = '80')&&(risk.info exists)$$time='2014-05-
04 18:50:00"-"2014-05-09 18:59:59(attachment exists) && (tcp.dstport = '80')
&& (risk.info exists) && time="2014-05-04 18:50:00"-"2014-05-09 18:50:59"
```

In addition, the time taken to load is displayed at the end of each meta key in the Values panel.

Time Banner

Just below the breadcrumb and debug information (if present), the time banner shows the time range used to create the chart. The following figure is an example of the time banner.

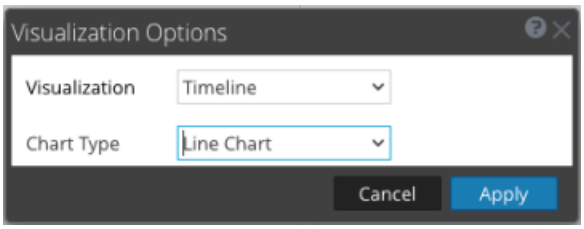


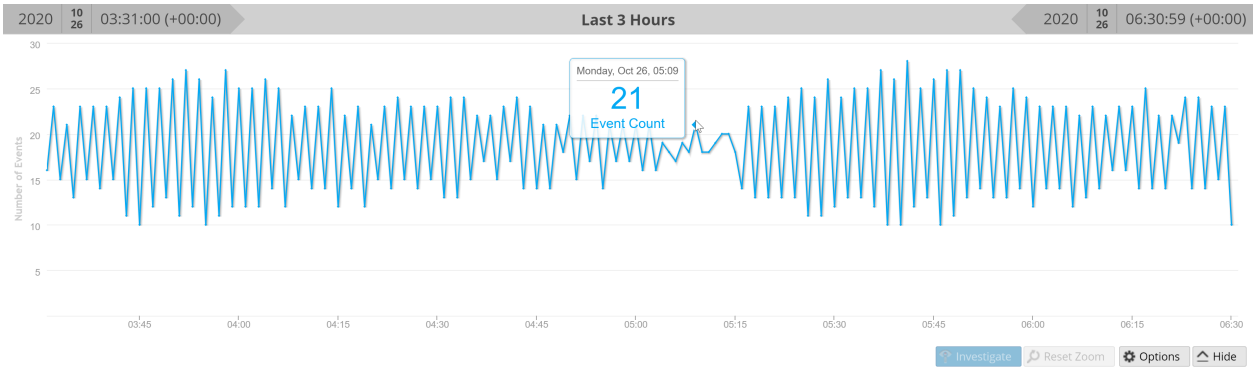
Visualizations

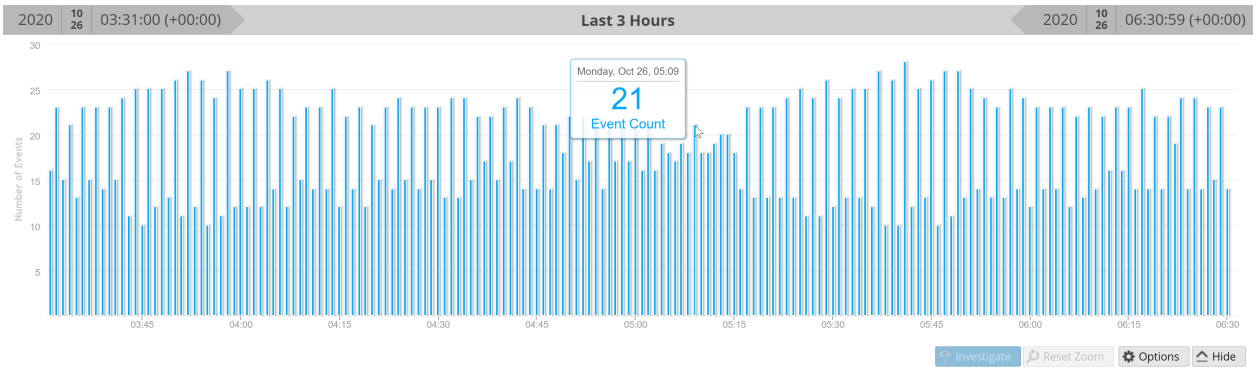
At the top of the Navigate view is a visualization of the current drill point. You can use this to drill into data from the Visualization panel (see Filter Results in the Navigate View). You can show or hide the visualization, and choose one of thevisualization options: Timeline or Coordinates. The Visualization opens initially to the last saved Visualization.

Timeline Chart

The timeline is the count of the number of events that occur at a specific instance. The timeline provides event counts so that you can see if the number of events increases drastically at a given point in time. The timeline displays activity for the specified service and time range as a line chart or a bar chart based on your choice in the Options menu. The second figure illustrates a line chart and third figure illustrates a bar chart.







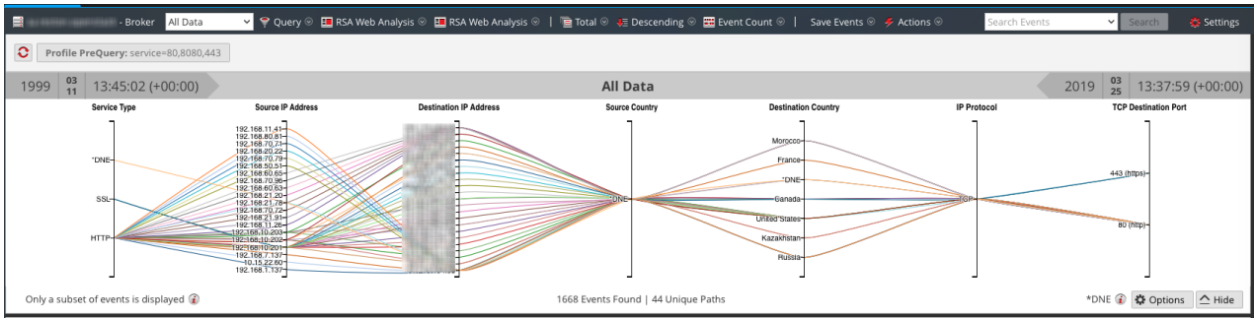
The timeline displays activity for the specified service and time range, as a line chart or a bar chart based on your choice in the Options menu.

Feature	Description
Number of Events (Timeline)	The Y axis of the chart based on thousands of events.
Time Line (Timeline)	The X axis of the chart based on the time the events occurred.
Event point (Timeline)	If you want to explore a specific section, simply select the range from the chart. The new time range will be reflected in the chart.
Investigate (Timeline)	Displays the meta values for the selected subset.
Reset Zoom (Timeline)	To return to the original time range, click Reset Zoom.

Feature	Description
Options	Displays the Visualization Options dialog. Data points can be displayed as a Line chart (default), a Bar chart, or Coordinates chart. When a chart type is select, the relevant options are displayed.
Hide	Collapses the chart.

Parallel Coordinates Chart

The Parallel Coordinates chart is one of the choices in the Options menu for visualizing the current drill point. With Coordinates selected in the Visualization Options dialog, you can select the meta data to be displayed (see Visualize Metadata as Parallel Coordinates). An easy way to view a useful Parallel Coordinates chart is to choose a profile group as shown in the following figure.



Feature	Description
Axes	Each axis is a meta key. The number of meta keys affects the load time for the chart. All meta keys are loaded, but it there the number of events per meta key is limited.
Lines	Lines represent events and they connect values on the axes to show the correlation between multiple meta keys.
Options	Displays the Visualization Options dialog. Data points can be displayed as a Line chart (default), a Bar chart, or Coordinates chart. When a chart type is select, the relevant options are displayed.
Only a subset of events is displayed.	This message is a notification that not all events in the values panel are drawn in the chart. Removing axes or filtering the data in the Values panel can help to display all events.
Events Found | Unique Paths	Displays the total number of events charted versus the number of unique paths charted. Setting the All Meta Keys Must Exist in an Event option redraws the chart so that it is more targeted and legible.
DNE	Indicates that there is no values for this meta key in the event.

In the Visualization Options dialog for Coordinates, you can select the meta keys to chart.

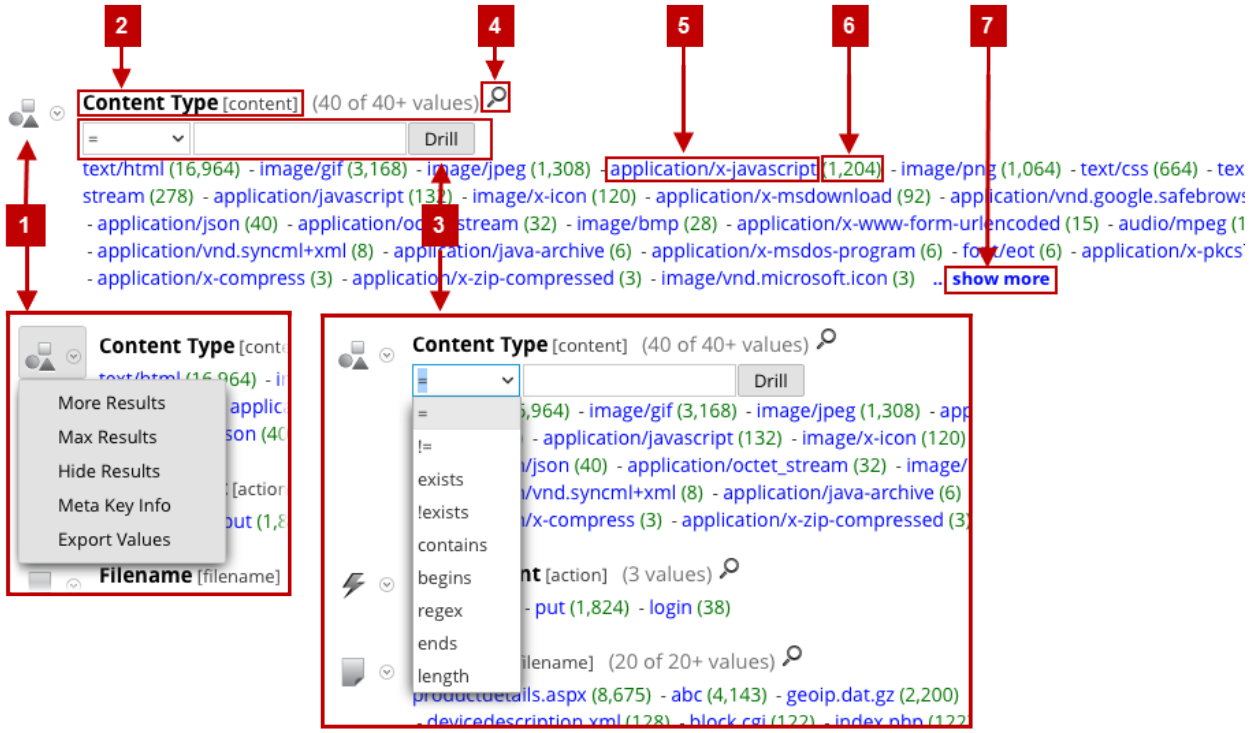
Feature	Description
Visualization selection	Displays a drop-down list of visualization types: Timeline and Coordinates
All Meta Keys Must Exist in an Event	Limits the data represented in the visualization to only those events that include all selected meta keys. This can result in a cleaner, more targeted visualization.
	Displays the Add Keys to Parallel Coordinates Visualization dialog so that you can add axes to the visualization. This is useful if you are looking for relationships between the default meta keys and some additional ones.
	Deletes the selected keys so that they do not appear as axes in the visualization. This can help to make the visualization less cluttered and allow for more data points to be included in the visualization.
	Reverts to the default meta keys for visualization, which consist of all meta keys in the current drill point.
	Controls the display of additional information about the number of selected axes versus the recommended count. This helps to make you aware of possible performance improvements by removing axes.
Axes	Lists the meta keys selected as axes in the visualization.
Cancel	Cancels any changes made to the visualization options.
Apply	Saves the changes made to the visualization options and applies to the current visualization.

In the Add Keys to Parallel Coordinates Visualization dialog, you can select the meta keys or meta groups to use as axes the Parallel Coordinates visualization.

Feature	Description
Visualization selection	Select Keys: Two options for selecting meta keys are: • From Default Meta Keys • From Meta Groups Each option offers a drop-down list from which to select.
With the Selected Meta Keys...	The options for the method of adding meta keys allow you to: • Replace the current list of keys • Append to the current list of keys • Insert at beginning of the current list of keys
Cancel	Closes the dialog and does not add any keys.
Add	Closes the dialog and adds the selected keys as specified.

Values Panel

The major feature of the Navigate view is the Values panel, which presents meta keys and meta values found in the service being investigated. Procedures for analyzing data in the Values panel are provided in Filter Results in the Navigate View.



Note: Title, values, and counts for non-indexed meta keys are not drillable; the Values and counts are shown in black.

1 The meta keys in the Values panel have drop-down menus, which offer actions that can apply to that meta key. You can use these to change the way the results for the meta key are displayed in the current view. Changes made to meta keys are displayed in the current view and persist until you refresh the page or select a new service in the Navigate view toolbar. See Drill into Data in the Values Panel

Refresh reverts to the current view of meta keys as defined in the Manage Default Meta Keys dialog (see Manage and Apply Default Meta Keys in an Investigation). If you have never made modifications in the Manage Default Meta Keys dialog, NetWitness, a refresh restores the default meta keys from the core service.

- More Results
- Max Results
- Hide Results
- Meta Key Info

	• Export Value
2	The name of the meta key for which values are listed. In Version 11.3 and later, the user friendly name of the meta key is displayed with the index file name of the meta key following in brackets. For example **Content Type [content]** gives the user friendly name of the `content` meta key with the index file name in parentheses. For meta groups, the name of the group is given in plain English with the meta group name following in parentheses. This is an example of a meta group name as it would appear in the Values panel: **All User Keys [users.all]**.
3 and 4	Clicking on an indexed meta key opens the Search dialog in which you can enter a filter for the current meta key. The search function is not available for non-indexed meta keys, and is based on the actual meta value rather than the alias. Drilling in the Search dialog using aliases is not supported. NOTE: Check with your administrator to obtain a list of aliases used for a meta key in Investigation. When an alias is used, this search dialog does not provide results. Instead, you must query the meta key using the Right-click query capability or the Query dialog.
5	The meta value associated with the found meta key. These are listed in order by meta value name or by the count of events in which the meta value was found, according to your preference.
6	The number of events that include the meta value.
7	The number or values rendered is specified by the Render Threads value in the Investigation Preference settings. In the example above, the meta key is **Content Type**, and 40 of 40+ values are currently displayed. You can display additional values by clicking **...show more**. The number of instances found for a particular meta in the session.

Values Panel Loading Behavior

The default view is for the last 3 hours of collection, using the default meta keys and non-indexed meta keys closed. The meta keys within the meta groups are displayed in the order that NetWitness queries the keys. As the data loads into the Values panel, NetWitness is optimized to show partial results, loading progress, and service status as the data loads.

The loading behavior is determined by several configuration settings. The highest level settings are configured by the administrator for each user. These are:

- The maximum amount of time allowed for this user to run a query (Query Timeout).
- The limit at which NetWitness stops counting the number of meta values in a session (Session Threshold). If a threshold is set for a session, the Navigate view shows that the threshold was reached and the percentage of results loaded. Any session that does not show a percentage is accurate and was processed to completion. If there is a percentage, that reflects how much processing was completed. The percentage displayed is estimated by extrapolating from the value at the time processing finished, considering the amount of work remaining. Larger percentages are generally more accurate because they require less extrapolating
- The limit at which NetWitness stops counting the number of meta values in a session (Session Threshold). If a threshold is set for a session, the Navigate view shows that the threshold was reached and the percentage of query time used to reach the threshold.

Note: The values for non-indexed meta keys take longer to load in the Values panel. To optimize loading, NetWitness does not open non-indexed meta keys by default. Refer to Manage and Apply Default Meta Keys in an Investigation for a detailed description of non-indexed meta keys in Investigation.

When you have launched an investigation of a service, NetWitness displays results in the Values panel.

1. NetWitness loads meta keys and meta values in the Values panel. For each meta key load, the stages of load are:
   a. **Waiting to Be Loaded** or **Closed**. If Closed, no data for that key is loaded.
   b. **Loading**
      i. **Loading progress**: NetWitness is receiving and displaying progress messages.
      ii. **Partial results**: NetWitness is receiving values messages and partial results are displayed in the Values panel.
   c. **Load Complete**: All results are finished loading.
2. As each meta key load is completed, and final values are displayed, the next meta key is started. The number or values rendered for each meta key is specified by the Render Threads value in the Investigation Preference settings. Loading continues until all keys to be loaded have finished.
3. If **Show Debug Information** is active and the service you are navigating is a 10.4 or later Broker, NetWitness displays load time information beneath the values for each meta key and displays additional load details for the aggregated services. NetWitness also displays the debug information beneath the breadcrumb.

Iterative results

Iterative results provide feedback on the status of queries within the interfaces to provide additional context for how long the data load will take and if any service data is missing. For example, if you are querying a Broker that is aggregating from two Concentrators, NetWitness starts displaying the results from the first Concentrator as soon as it is available, even if the second Concentrator is still waiting for results.

Iterative results also include a notification that service data is missing because the service is unreachable.

Partial results

When partial values from the Core service are returned but not completed, a message at the end of the meta key listing shows the progress of values loaded. For example, Currently looking at 38 ip.src values 71% indicates that loading of values for the meta key is 71% complete.

Debug Information

If the Show Debug Information setting is in effect, a field at the end of the values displays the status for the different systems against which you are querying within NetWitness. For example, when you are querying against a 10.4 broker pulling from multiple concentrators, NetWitness displays the status of the query on each of the Concentrators, which provides insight into the relative speed of data loading from each of the Concentrators. Each service that participated in the query is listed with the total elapsed time for the query.

Each service that participated in the query is listed with the total elasped time for the query. In the example above, two services returned in 3.207 seconds, localhost:50005 took 2 seconds to return the results. In addition, the where clause of the query is displayed below the breadcrumb. You can copy this syntax directly into an application rule or Reporting where clause of a rule.

Load Complete

For each meta key, there is a list of values (blue text) and counts (green text) found in the current drill point. When you click a value to drill down into a subset of the currently selected data, the display is updated and the new drill point is recorded in the breadcrumb. You can specify the sorting and quantification methods for the values list using the option in the toolbar.

Query Dialog

In the Navigate view or Legacy Events view, you can create a query rather than clicking through the meta keys and values to drill down into the meta data. The dialogs for creating a query offer syntax help with drop-down lists of applicable meta keys and operators. To access this dialog in the **Navigate** or **Legacy Events** view toolbar, select **Query**.

What do you want to do?

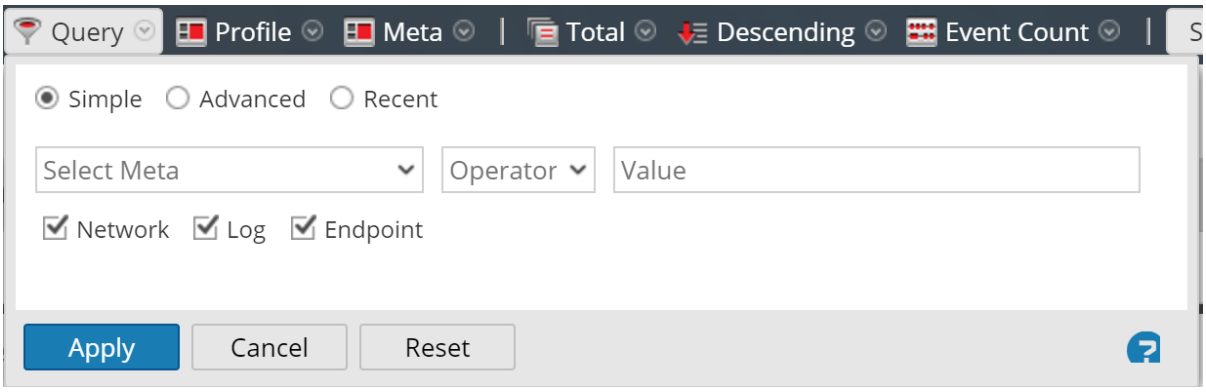
User Role	I want to ...	Show me how
Incident Responder or Threat Hunter	review detections and signals seen in my environment	*NetWitness Platform Getting Started Guide*
Incident Responder	review critical incidents or alerts	*NetWitness Respond User Guide*
Threat Hunter	query a service, metadata, and time range*	Begin an Investigation in the Events View Begin an Investigation in the Navigate or Legacy Events View
Threat Hunter	view metadata	Filter Results in the Navigate View Drill into Metadata in the Events View (BETA)
Threat Hunter	view sequential events	Filter Results in the Events View Filter Results in the Legacy Events View
Threat Hunter	reconstruct and analyze an event	Examine Event Details in the Events View Reconstruct an Event in the Legacy Events View
Threat Hunter	examine files and associated hosts	Download Data in the Events View Export or Print a Drill Point in the Navigate View Export Events in the Legacy Events View
Threat Hunter	perform lookups	Look Up Additional Context for Results Launch a Lookup of a Meta Key
Threat Hunter	create an incident or add to an incident	Add Events to an Incident in the Legacy Events View Add Events to an Incident in the Events View
Threat Hunter	add a meta value to a Context Hub list	Look Up Additional Context for Results

*You can perform this task in the current view.

Related Topics

- How NetWitness Investigate Works
- Navigate View
- Legacy Events View

Quick Look

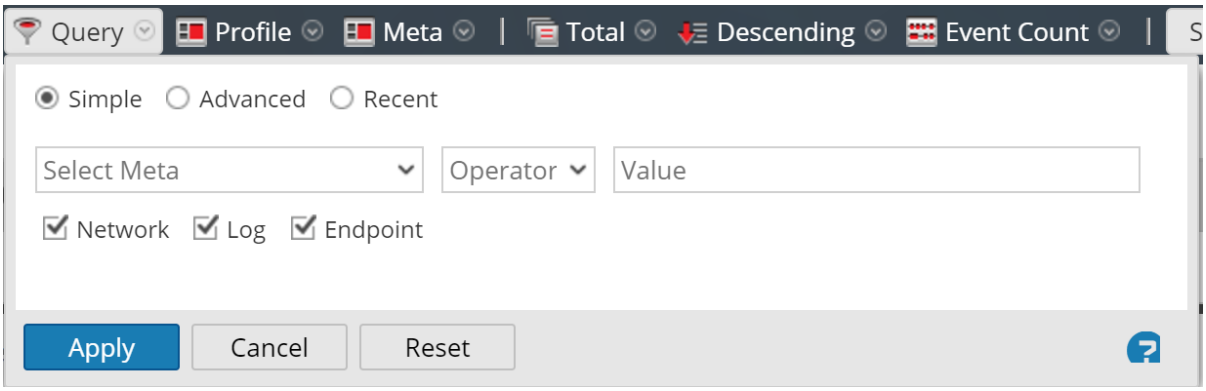


The Query dialog has three views:

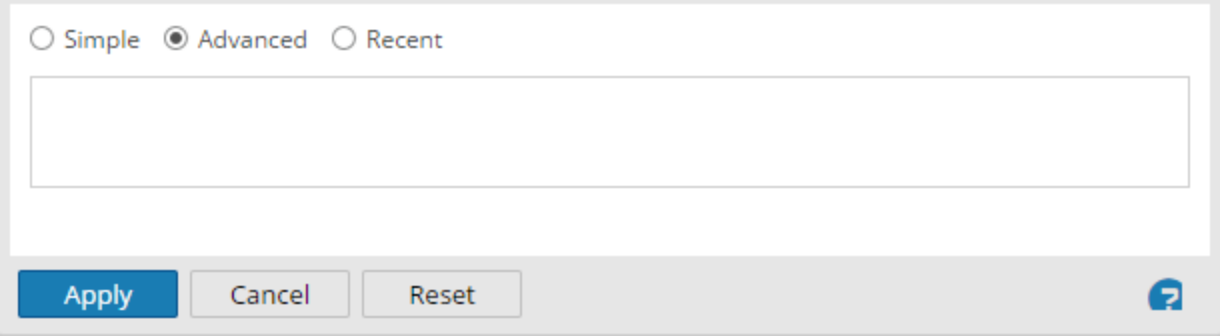
- Simple
- Advanced
- Recent

In the Simple view, you can create a query using the options displayed in the dialog. In the Advanced view, you can create a query without guidance. In the Recent view, you can select a query from a drop-down list of recent queries.

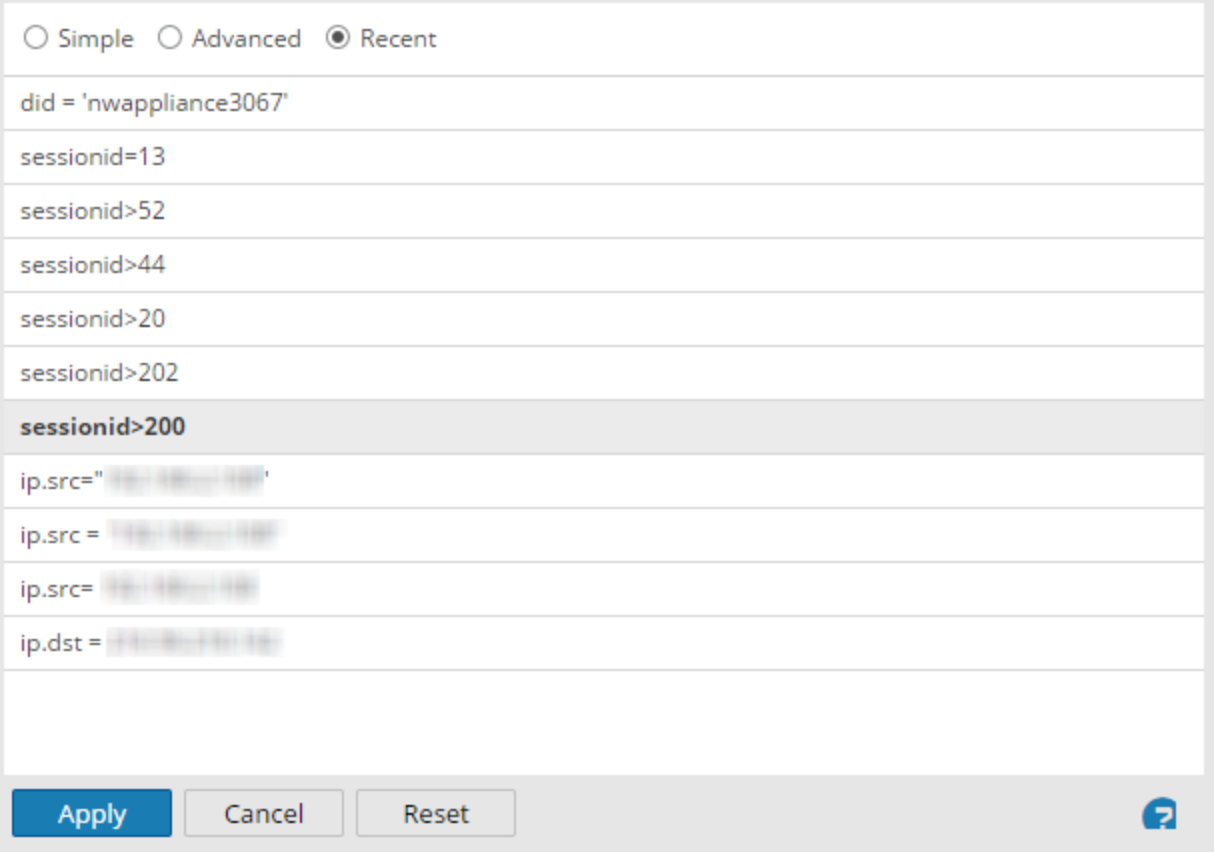
Simple View



Advanced View



Recent View



The following table describes features of the Query dialogs.

Feature	Description
Select Meta	Displays a drop-down list of meta groups.
Operator	Displays a drop-down list of operators (=,NetWitness!=,NetWitnessexists,NetWitness!exists)

Feature	Description
Value	Allows you to enter a value to complete the query.
Network	Limits the query to packets if Log is not selected.
Log	Limits the query to logs if Network is not selected.
Query box	Allows you to enter a query in the Advanced view. When you begin typing, a drop-down list of available meta keys for the service is displayed, then a drop-down of operators is displayed as you type. If the expression currently entered in the query box is invalid, a warning appears near the box. When the query is valid, the warning is removed.
Query list	Allows you to select a query from a list of recent queries in the Recent view. Double-clicking a query automatically applies it.
Apply	Applies the new query to the current Investigation view.
Cancel	Closes the dialog without applying changes.
Reset	Resets all fields.

Query Profiles Dialogs

Query profiles offer a quick and easy way to define a meta group, column group, and a limiting filter (pre-query condition) that you can apply in the Navigate view, the Events view, and the Legacy Events view (see Use Query Profiles to Encapsulate Common Areas for Investigation). The same query profiles are shared between all views, and they are available in the Springboard (Version 11.5) for use in panels. Private query profiles created in the Events view are only available in the Events view for the analyst who created them.

Each query profile specifies a meta group, column group, and sometimes includes a pre-query condition appropriate for the type of investigation.

In a query profile:

- The meta group defines the meta keys that are queried (see Use Meta Groups to Focus on Relevant Meta Keys).
- The column group defines which meta keys from the meta group are displayed as columns in the Events list. (see Use Columns and Column Groups in the Events List).
- When the query profile is in effect, the optional pre-query conditions add a limiting filter in the query bar. You can edit or delete the limiting filter and then create additional filters for your query (see Filter Results in the Events View)

You can manage profiles in the Manage Profiles dialog, the Create Query Profile dialog and the Query Profile Details dialog.

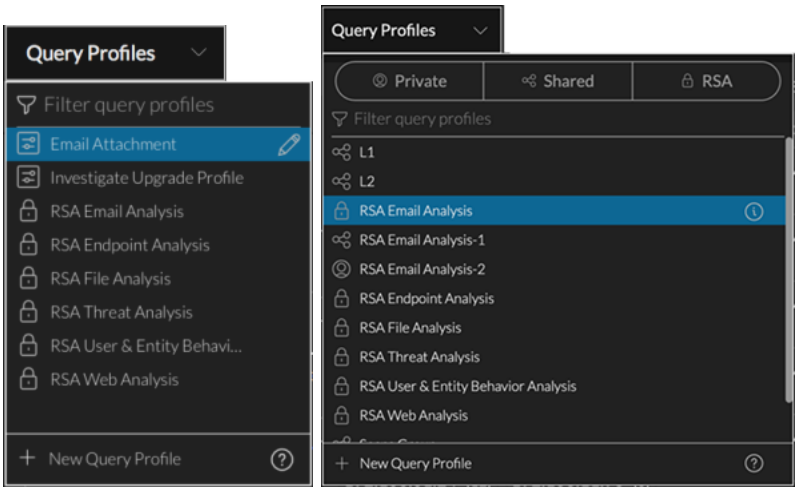
- The Manage Profiles dialog is for the Navigate view, the Legacy Events view (Version 11.4 and later) , and the Events view (Version 11.3 and earlier). To access this dialog, select **Profile > Manage Profiles** in the **Navigate** or **Legacy Events** view toolbar.
- The Create Query Profile dialog is for the 11.4 and later Events view. To access this dialog, select **Query Profiles > New Query Profile** in the **Events** view query bar.
- The Query Profile Details dialog is for the 11.4 and later Events view. To access this dialog, select **Query Profiles** in the **Events** view query bar, then click the edit icon () next to a custom profile name.

Related Topics

- How NetWitness Investigate Works
- Use Query Profiles to Encapsulate Common Areas for Investigation
- Navigate View
- Events View
- Legacy Events View

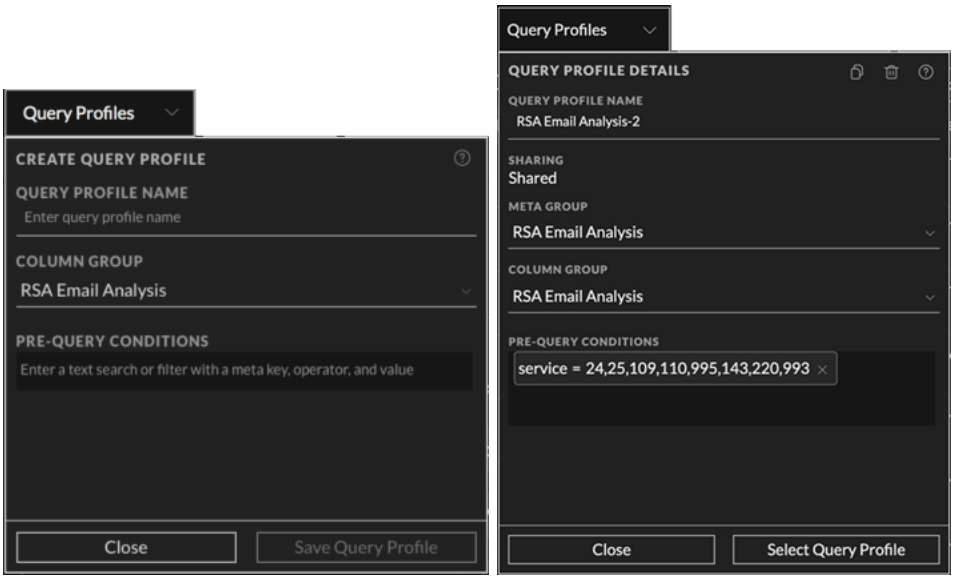
Quick Look - Query Profile Menu, Create Query Profile Dialog, and Query Profile Details Dialog

This section introduces the Query Profile menu, Query Profile dialog, and the Query Profile Details dialog. The following figure is an example of the Query Profiles menu and the table describes the options. The example on the left has built-in profile highlighted so that the information icon is visible. The Version 11.4 menu is on the left and the Version 11.5 menu is on the right.



Feature	Description
Visibility Options	Control the types of query profiles that are visible in the list. You can use any combination of the visibility options: **Private**, **Shared**, or **RSA** (blue = selected, black = not selected). Initially none of the buttons are selected and all profile types are visible. This is the same result as if all three buttons are selected. The visibility options work together with text in the Filter Query Profiles field. If the visibility option is hiding built-in profiles (which include "RSA" in the name) and you search for a name that contains "RSA," the list is empty. Private = display private groups that only you can manage Shared = display shared groups that anyone in your organization can manage RSA = display built-in groups that only RSA can manage
Filter Query Profiles	Filters the list of query profiles as you type text so that only profile names that contain that text are displayed.
Query Profile List	The list of profiles consists of custom and built-in profiles, which are distinguished by the icons that precede the name. In the example, RSA Email Analysis-1 and RSA Email Analysis-2 are custom profiles. The RSA Email Analysis is a built-in profile.
New Query Profile	Displays the Create Query Profile dialog, where you can create a custom profile.

The Create Query Profile dialog, shown in the figure on the left, allows you to define a custom profile. The figure on the right illustrates the Query Profile Details dialog, in which you can edit a custom profile. The table describes the fields and options in the dialogs.

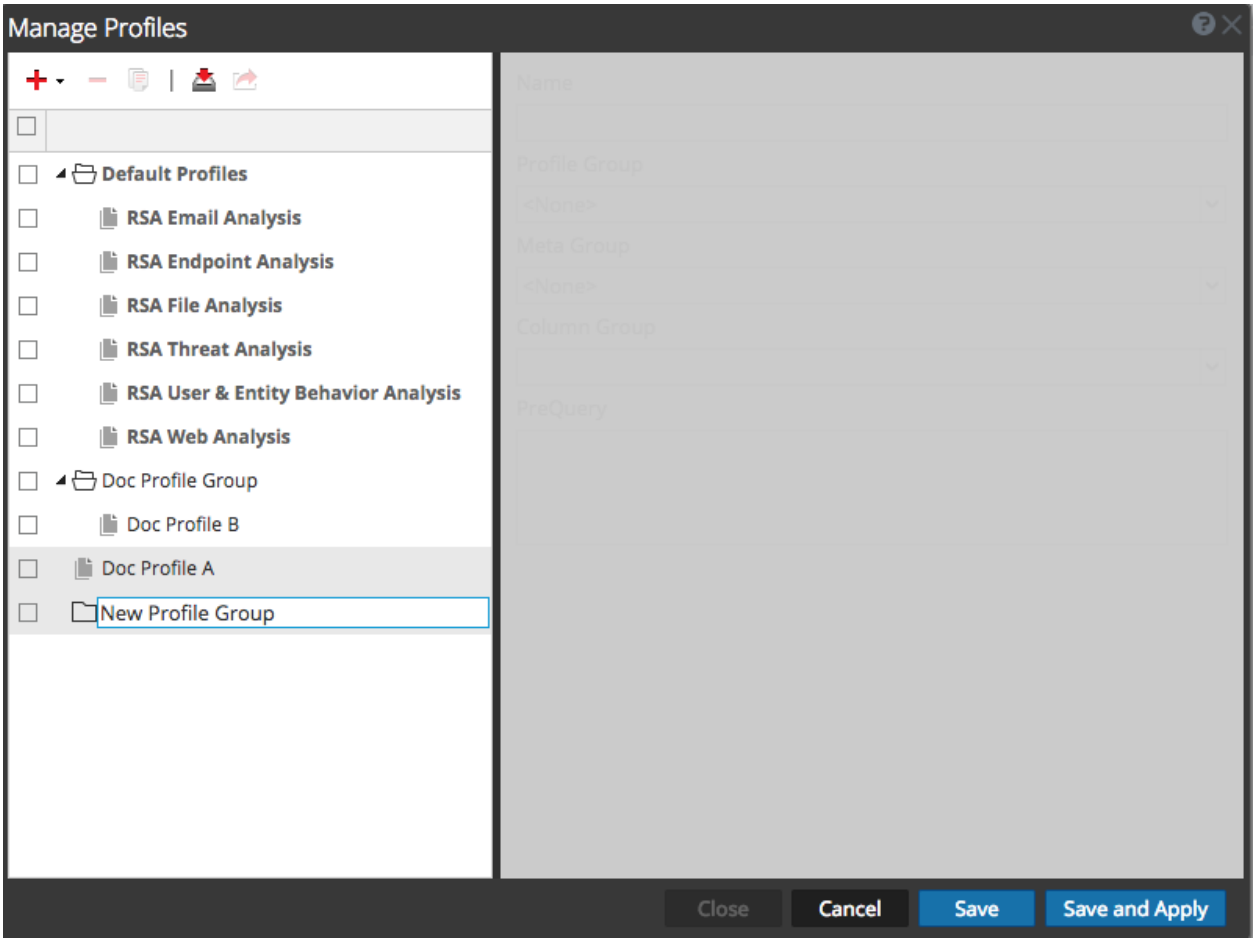


Feature	Description
	Creates a clone of the meta group so that you can edit a copy. This is useful if you want your own copy of a built-in group, a shared copy of a private group, or a private copy of a shared group.
	Deletes the custom profile in the Query Profile Details dialog. This action is irreversible and applies globally; the profile is no longer available to anyone who is using the profiles on this service.
Query Profile Name	Displays the name of the profile. The name must be unique and contain fewer than 64 characters. You can edit the name in a custom profile.
Column Group	Displays a drop-down menu listing available column groups, with the currently selected column group from the Events list already selected. You can change the column group in a custom profile.
Pre-Query Conditions	Defines a limiting filter for results in the Events view. If you had a query active in the query bar when you began to create the new profile, the active query is added to the pre-Query field. In a custom profile, you can delete the prepopulated pre-query condition and type additional text for a text search or additional filters in the Pre-Query Conditions field. This is an example of a pre-query condition: `'service=80,25,110'`.
Close button	Closes the dialog.
Save Query Profile	For the Create Query Profile dialog only, saves the new profile.

Feature	Description
Reset	For the Query Profile Details dialog only, reverts the edited profile to the last saved state.
Update Query Profile	For the Query Profile Details dialog only, applies changes to an edited profile.
Select Query Profile	Applies the query profile.

Quick Look - Manage Profiles Dialog

This is an example of the Manage Profiles dialog showing several profile groups.



The Profile panel on the left side of the dialog displays available profiles and allows you to add, delete, import, and export profiles. The following table describes the fields in the Profile panel.

Field	Description
	Adds a new profile using the Settings panel on the right side of the Manage Profiles dialog.
	Deletes the selected profile. A confirmation dialog is displayed before the profile is deleted.
	Creates a copy of the selected profile.
	Displays the Profile Import dialog, where you can upload a file.
	Exports the selected profile to your computer.
Profile Name	Lists all profile names.

The Settings panel on the right side of the dialog offers options to configure profiles. It can only be used when one profile is selected. The following table describes the fields in the Settings panel.

Feature	Description
Name	Displays the name of the profile.
Meta Group	Displays a drop-down menu listing available meta groups.
Column Group	Displays a drop-down menu listing available column groups. The OOTB column groups and these three groups are available by default: • List View • Detail View • Log View
PreQuery	Defines a limiting query for filtering Investigate results. This query is used when the associated profile is activated and the preQuery applies to any queries used in the Navigate and Events views. This is an example of a preQuery: `'service=80,25,110'`.

The following table describes the buttons.

Field	Description
Close	Closes the dialog.
Cancel	Cancels all changes.
Save	Saves all changes.
Save and Apply	Saves and applies all changes immediately.

Settings Dialogs for Investigate Views

NetWitnessVersion 11.0 has two settings dialogs, one for the Navigate view and one for the Legacy Events view. With the addition of the settings dialog for the Events view in Version 11.1, Investigate has three settings dialogs.

The settings in these dialogs are a subset of the Investigation settings made in the Profiles > Preferences panel > Investigation. Analysts can save time by editing these settings within the Investigate view. If you change a setting here, the same setting is changed in the Profiles view, and if you change a setting in the Profiles view, the same setting is changed here.

To access this dialog, go to the **Navigate** or **Legacy Events** view, and select the **Settings** option in the toolbar.

The settings in the Events view have no corresponding settings in the Profiles > Preferences panel.

Related Topics

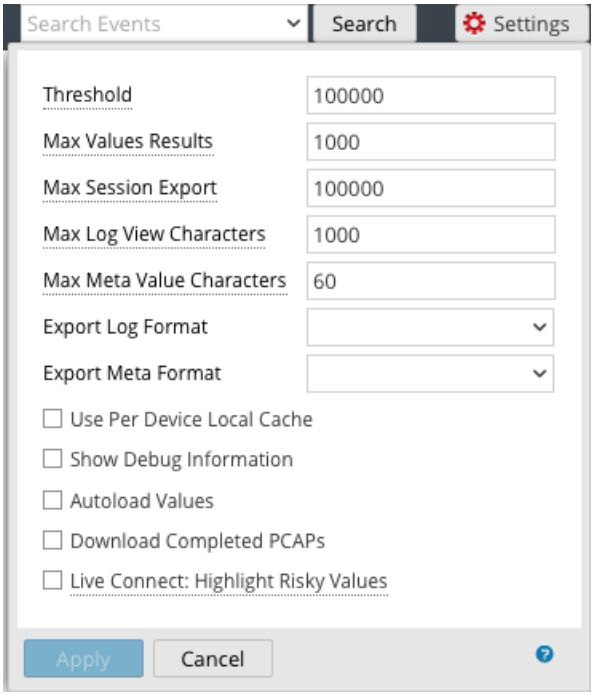
- How NetWitness Investigate Works

Quick Look

This is a quick look at the settings dialog for the Navigate view, Legacy Events view, and Events view.

Navigate View Settings Dialog

The following figure illustrates the Navigate view Settings dialog. The settings that influence performance when loading values in the Values panel have default values based on common usage, and individual analysts can adjust these settings for their own investigations. The following table describes the features.

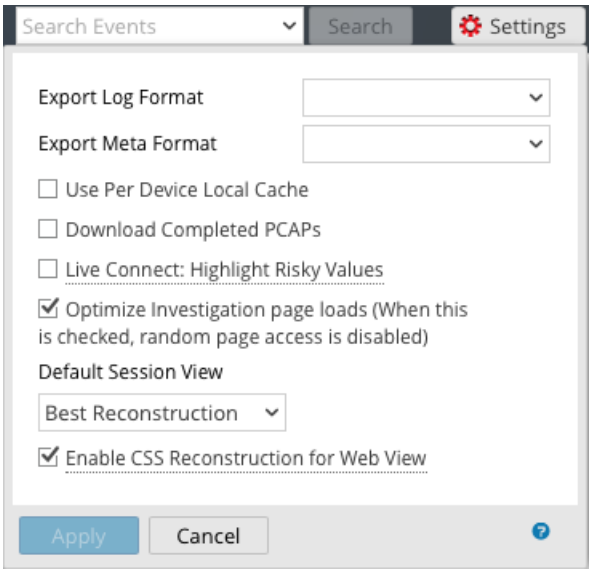


Feature	Description
Threshold	Sets the threshold for the maximum number of sessions loaded for a meta key value in the Values panel. A higher threshold allows accurate counts for a value, and also causes longer load times. The default value is **100000**.
Max Values Results	Sets the maximum number of values to load in the Navigate view when the Max Results option is selected in the Meta Key menu for an open meta key. The default value is **1000**.
Max Session Export	Sets the maximum number of sessions that can be exported. The default value is **100000**.
Export Log Format	Sets the file format of exported logs. There are four formats available: • **Text**: Raw log format. • **SML**: Structured Markup Language format. • **CSV**: Comma-separated values (CSV) format. • **JSON**: The JavaScript Object Notation (JSON) format.

Feature	Description
Export Meta Format	Sets the file format of exported meta values. There are four formats available: • **Text**: Raw log format. • **SML**: Structured Markup Language format. • **CSV**: Comma-separated values (CSV) format. • **JSON**: The JavaScript Object Notation (JSON) format. **Note:** If you upgrade to version 11.5.2, the **Export Meta Format** preference is not retained and is reset to blank. You must re-configure this value after you upgrade to version 11.5.2.
Use Per Device Local Cache	When the checkbox is cleared, Investigate sends a fresh query to the database rather than displaying cached data in the Investigate views after the initial load. If the checkbox is selected, Investigate uses the data from local cache.
Show Debug Information	This option controls the display of the `where` clause beneath the breadcrumb in the Navigate view and the elapsed load time for each aggregated service on a Broker. When the checkbox is selected, the debug information is displayed. The default value is **Off** (checkbox cleared).
Autoload Values	This option controls automatic loading of values for the selected service in the Navigate view. When the checkbox is selected, values are automatically loaded when you select a service to investigate. When the checkbox is cleared, Investigate displays a **Load Values** button, allowing the opportunity to modify options. The default value is **Off**.
Download Completed PCAPs	This setting automates the downloading of extracted PCAPs in Investigate so that you do not have to manually download and open extracted PCAP files in an application, such as Wireshark, that can handle viewing data in a PCAP form. When the checkbox is selected, the option is enabled. The default setting is disabled (checkbox is cleared).
Live Connect: Highlight Risky IPs	If the checkbox for this option is cleared, all the meta values that have context available in Live Connect are highlighted in the Navigate view Values panel. If the checkbox is selected, among the values that have context in Live Connect, only those values deemed Risky/Suspicious/Unsafe by the community are highlighted. By default this option is disabled (checkbox cleared).
Apply	Applies the settings immediately and they are visible the next time you load values. The same changes are also applied in the Profiles view.
Cancel	Cancels the editing operation and closes the dialog, leaving the settings unchanged.

Legacy Events View Settings Dialog

The following figure is an example of the Settings dialog for the Legacy Events view, and the table describes the features.

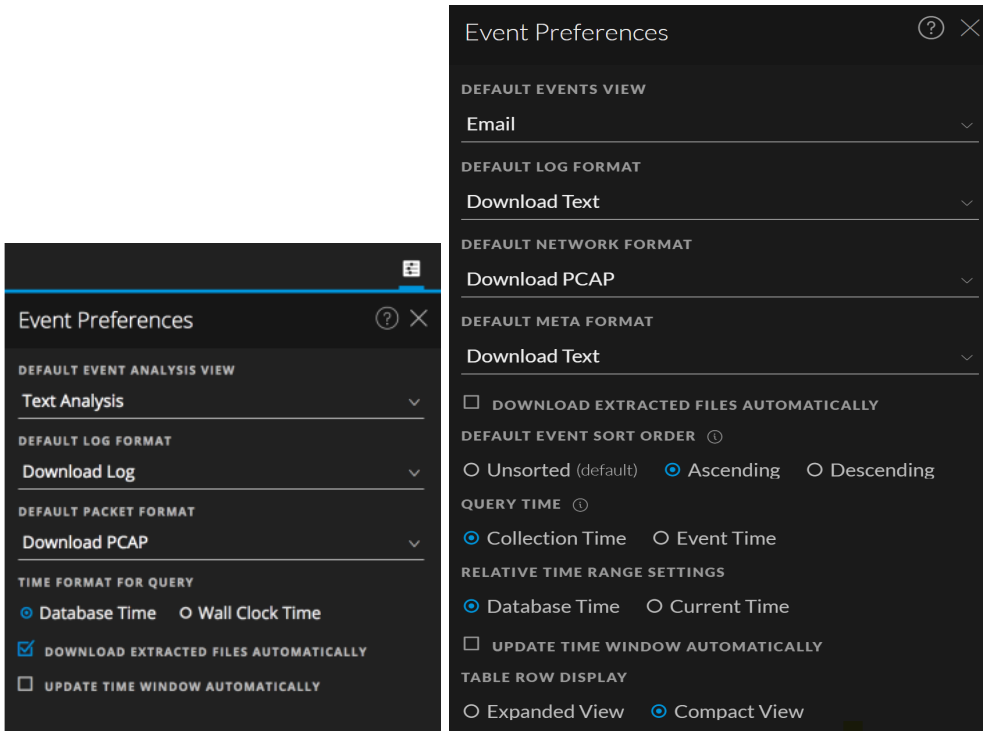


Feature	Description
Export Log Format	Sets the file format of exported logs. There are four formats available: • **Text**: Raw log format. • **SML**: Structured Markup Language format. • **CSV**: Comma-separated values (CSV) format. • **JSON**: The JavaScript Object Notation (JSON) format.
Export Meta Format	Sets the file format of exported meta values. There are four formats available: • **Text**: Raw log format. • **SML**: Structured Markup Language format. • **CSV**: Comma-separated values (CSV) format. • **JSON**: The JavaScript Object Notation (JSON) format.
Download Completed PCAPs	This setting automates the downloading of extracted PCAPs in Investigate so that you do not have to manually download and open extracted PCAP files in an application, such as Wireshark, that can handle viewing data in a PCAP form.
Live Connect: Highlight Risky IPs	When the checkbox is selected, Investigate uses a filter to fetch only IP addresses that are considered as risky by RSA community. When the checkbox is cleared, NetWitness displays all IP addresses. By default, this option is disabled (checkbox cleared).

Feature	Description
Optimize Investigation page loads	Sets a paging option. When optimized, results are returned as quickly as possible, sacrificing the original ability to go to a specific page in the event list. Clearing this checkbox changes the Events list pagination to allow you to go to a specific page in the list (or to the last page). The default value is enabled (checkbox selected).
Append Events in Event Panel	This option affects paging in the Legacy Events panel and in prior releases was located in the Navigate view settings dialog. When the checkbox is selected, the next group of events is appended to the already displayed events. When cleared, the previous page of events is replaced by the next page. The default value is Off (checkbox cleared).
Default Session View	Selects the default reconstruction type for the initial reconstruction in the Events view. The default value is **Best Reconstruction** in which events are reconstructed using the reconstruction method most appropriate to the event.
Enable CSS Reconstruction for Web View	This setting controls how web content reconstruction is performed. If enabled, the web reconstruction includes cascaded style sheet (CSS) styles and images so that its appearance matches the original view in a web browser. This includes scanning and reconstructing related events, and searching for style sheets and images used in the target event. The option is enabled by default. Clear the checkbox to disable this option if there are problems viewing specific websites.
Apply	Applies the settings immediately and they are visible the next time you view events. The same changes are also applied in the Profiles view.
Cancel	Cancels the editing operation and closes the dialog, leaving the settings unchanged.

Events View Preferences Dialog

Beginning with Version 11.1, the Events view has user preferences that you can configure in the Events view > Event Preferences dialog. These settings persist so that they are applied each time you log in and go to the Events view. The following figures are examples of the dialog for Version 11.3 and Version 11.6. The table below describes the options.



Feature	Description
Default Events View	Selects the default event analysis view that is displayed every time you open the Events view. For example, if you select File, the File Analysis panel is highlighted and displayed every time you investigate an event in the Events view. These are following options: • **Text** : View and analyze the raw text payload of an event. • **Packet** : View and interactively analyze the packets and payload of an event. • **File** : View a list of files and download one or more files in an event.
Default Log Format	Selects the default format for downloading logs: • **Download Log** or **Download Text**: Raw log (log) format. • **Download CSV**: Comma-separated values (CSV) format. • **Download XML**: The Extensible Markup Language (XML) format. • **Download JSON**: The JavaScript Object Notation (JSON) format.

Feature	Description
Default Packet Format or Default Network Format	Selects the default format for downloading packets. • **Download PCAP**: To download the entire event as a packet capture (*.pcap) file. • **Download All Payloads** or **Download Payloads**: To download the payload as a *.payload file. • **Download Request Payload**: To download the request payload as a *.payload1 file. • **Download Response Payload**: To download the response payload as a *.payload2 file.
Default Meta Format	Selects the default format for downloading metadata: • **Download CSV**: Comma-separated values (CSV) format. • **Download JSON**: The JavaScript Object Notation (JSON) format. • **Download Text**: Plain text format. • **Download TSV**: Tab-separated values (TSV) format.
Time Format for Query	The Events view can display results based on the database time or the current clock time. **Note:** (Version 11.6) **Current Time** is the default for **Relative Time Range Settings**. In previous versions, **Database Time** was the default value. Make a note that this may cause time range mismatch between **Events View** (using **Current Time** as default) and **Navigate View** (using **Database Time** as default). This change does not affect the existing users and is applicable only to the new users. When **Database Time** is selected, the start and end time for a query is based on the time that the event was captured (collection time). When **Current Time** (labeled **Wall Clock Time** in Version 11.3 and earlier) is selected, the query is executed using the end time based on the current browser time; the start time is calculated based on that end time and the time range.

Feature	Description
Event Sort Order (Version 11.4 and Later)	Sets the sort sequence by collection time for the events listed in the Events panel. If results exceed the events limit, not all events can be loaded. The portion of returned events loaded in the Events panel matches the sort order preference: the oldest portion of events is loaded when Ascending order is selected, and newest portion of events is loaded when Descending order is selected. A change to this setting becomes effective the next time you submit a query. **Unsorted**: Default sorting method for Version 11.4.1. To list events as processed by the Core services. Unsorted is faster because it streams back the events as soon as a match is found versus waiting for all Core services to respond and then displaying them in the chosen order. **Ascending**: Default sorting method for Version 11.4. To put the events with the earliest collection time first in the list. **Descending**: To put the events with the latest collection time first in the list. When investigating logs, you may want to change the sort sequence to latest collection time first.
Download Extracted Files Automatically	Enables the automatic download of files if they are in the selected default format in the **Default Log Format** and **Default Packet format** fields set in the Event Preferences dialog. Select the checkbox to enable downloading the selected format automatically to local file system. Otherwise, the download job goes to the job queue, and you can download it manually.
Update Time Window Automatically	(Version 11.3 and later) Enables automatic update of the time range window in the query bar when the service is polled (at one minute intervals) so that fresh results are sent. The default setting is disabled. When the checkbox is selected, as the time range is updated, the  (Submit Query) button is activated so you can click to get the fresh results. When the checkbox is cleared, the automatic update is disabled keeping the time range window in the breadcrumb synchronized with the current results.