



ArcSight ESM

Software Version: 7.9

ArcSight Console User's Guide

Document Release Date: November, 2025

Software Release Date: November, 2025

Legal Notices

Open Text Corporation

275 Frank Tompa Drive, Waterloo, Ontario, Canada, N2L 0A1

Copyright Notice

Copyright 2001-2025 Open Text.

The only warranties for products and services of Open Text and its affiliates and licensors ("Open Text") are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Open Text shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

Trademark Notices

"OpenText" and other Open Text trademarks and service marks are the property of Open Text or its affiliates. All other trademarks or service marks are the property of their respective owners.

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number
- Document Release Date, which changes each time the document is updated
- Software Release Date, which indicates the release date of this version of the software

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

<https://www.microfocus.com/support-and-services/documentation>

Support

Contact Information

Phone	A list of phone numbers is available on the Technical Support Page: https://softwaresupport.softwaregrp.com/support-contact-information
Support Web Site	https://softwaresupport.softwaregrp.com/
ArcSight Product Documentation	https://www.microfocus.com/documentation/arcsight/

Contents

Chapter 1: Getting Started	40
Starting the ArcSight Console	40
Quick Start Tools and Standard Content	41
Use Cases	41
Chapter 2: Working in the	43
Navigating	43
Navigator Panel Resource Tree	44
Batch Editing	46
Batch-Editing Cases or Connectors	46
Locking Case Groups	46
SmartConnector Reminders	46
Reconnecting to the Manager	47
Changing the Console Display	47
Changing User Preferences	48
Changing Your Password	48
Setting Default Editors and Viewers	49
Changing Global Options	49
Setting Dialog Options	51
Setting Grid Options for the Viewer Panel	52
Customizing the Default Selections for Active Lists	54
Setting Date and Time Formats	55
Setting Latitude and Longitude Options	56
Configuring Event Graphs	57
Setting Notification Popups	58
Managing Hot Keys	58
Adding Shortcuts for Frequently-Used Resources	59
Modifying a Custom Shortcut	62
Removing a Custom Shortcut	64
Activating a New Shortcut Schema	65
Sharing Custom Shortcut Schemas	65
Viewing	66
The Viewer Panel	66
Look-and-Feel	68
Inspecting and Editing	68

Overview of Inspect/Edit Features and Utilities	69
Searching for Fields in Event Inspector, Resource Editors, or CCE	70
Getting More Help	71
Controlling the	71
Using the Network Tools	73
Running a Tools Command	74
Adding or Editing a Tool	75
Staying Informed	76
Acknowledging Notifications	77
Checking the Status of the Distributed Correlation Cluster	77
Defining Message Lag Thresholds	78
Using Notes	79
License Tracking	80
License Tracking Notifications	81
Standard Reports for License Status Tracking	81
Using the File Menu	81
Using the Edit Menu	82
Using the View Menu	82
Using the Window Menu	84
Using the Tools Menu	84
Using the System Menu	85
Using the Help Menu	85
Using Right-Click Context Menus	86
Using the Advanced Selector While Editing Resources	88
Keyboard Shortcuts (Hot Keys)	89
Creating Shortcuts for Resources	91
Showing Recently Viewed Resources	91
Adding Resources to the Favorites List	92
Printing from the Console	93
Printing Navigation Tree Views of Resources	93
Printing Resource Definitions	93
Printing Grid Views	94
Printing Conditions Tree Summary	95
Using Column Flip Limit to Format Grid View Printouts	95

Saving and Sending Settings	97
Error and Warning Messages	98
Chapter 3: Managing Users and Groups	99
Managing User Groups	99
Managing Users	101
Creating or Editing a User	102
Resetting User Passwords	104
Moving or Linking a User	105
Deactivating and Reactivating a User	105
Deleting a User	106
Chapter 4: Managing Permissions	107
Editing Access Control Lists (ACLs)	107
Granting or Removing Resource Permissions	108
Granting or Removing Operations Permissions	109
Granting or Removing User Group Permissions	110
Adding or Removing Enforced Filters	113
Permissions for Sortable Field Sets	115
Sharing Resources	116
Controlling Who Has Permissions to Deploy Data Monitors	116
How Upgrades Affect Data Monitor Deploy Permissions	118
Deployment Permissions on Imported Data Monitors	118
Chapter 5: Modeling the Network	119
The Network Model	119
Assets	120
Automatically-Created Assets	121
Asset Aging and Model Confidence	123
Asset Ranges	123
Zones	124
Dynamic and Static Zones	125
Networks	126
Asset Model	126
Locations	126

Vulnerabilities	127
Asset Categories	127
Asset Categories Assigned to Assets, Asset Ranges, and Asset Groups	127
Asset Categories Assigned to Zones	128
Populating the Network Model with Assets	128
ArcSight Console-Based Methods	129
Manually, Using Network Modeling Resources	129
In a Batch Using the Network Modeling Wizard	130
SmartConnector-	130
Using the Asset Model Import FlexConnector	130
Automatically From a Vulnerability Scanner Report	131
ArcSight-Assisted Methods	131
As an Archive File From an Existing Configuration Database	132
Populating the Network Model Using the Wizard	132
Specifying CSV Column Types	133
Specify the Column Type Using a Header	133
Specifying Multiple Categories in one Category Column	134
Assign the Column Type in the Wizard	135
Zones CSV File Format	136
An Example of a Zones CSV File	137
Zones CSV File Format	137
An Example of a Zones CSV File	139
Assets CSV File Format	139
An Example of an Assets CSV File	141
Static Addressing in a Dynamic Zone	141
Asset Ranges CSV File Format	142
An Example of an Asset Ranges CSV File	143
Increasing the Number of Displayed Rows	143
Summary of Data to Import	144
Network Data Imported into ArcSight Manager	144
Working with Assets, Locations, Zones, Networks, Vulnerabilities, and Categories	145
Managing Assets	145
Asset Auto-Creation	147
Creating Assets from a Vulnerability Scan Report	148
Creating Assets from a Vulnerability Scan Report for Static Zones	148
Creating Assets from a Vulnerability Scan Report for Dynamic Zones	149
Creating Assets for SmartConnectors	149
Creating Assets for SmartConnectors in Static Zones	150

Creating Assets for SmartConnectors in Dynamic Zones	151
Creating Assets for Network Devices	151
Creating Assets for Network Devices in Static Zones	152
Creating Assets for Network Devices in Dynamic Zones	153
Asset Auto-Creation from Scanners in Dynamic Zones	153
Create Asset with IP Address or Host Name	154
Preserve Previous Assets	155
Asset Names	156
Changing the Default Naming Scheme	157
Selecting Assets in the Common Conditions Editor	158
Auto-Zoning an Asset	159
Auto-Zoning Imported Assets	160
Managing Asset Groups	161
Managing Vulnerabilities	162
Selecting Vulnerabilities in the Common Conditions Editor	164
Working with Vulnerable Assets	165
Managing Vulnerability Groups	166
Showing Affected Assets	167
Reporting on Output from Vulnerability Scanners	168
Reporting on Asset Vulnerabilities	168
Managing Zones	168
Managing Networks	170
Managing Asset Categories	171
Managing Locations	172
Managing Customers	172
 Chapter 6: Managing SmartConnectors	 175
Selecting and Setting SmartConnector Parameters	175
Configuring the SmartConnector	175
Connector Editor Tabs	176
Connector Tab Configuration Fields	177
Default Content Tab Configuration Fields	178
SmartConnector Processing Categories	190
SmartConnector Time Interval Options	191
Managing SmartConnector Filter Conditions	192
Adding SmartConnector Filter Conditions	192
Deleting SmartConnector Filter Conditions	193
Setting Special Severity Levels	193

Sending Model Mappings to SmartConnectors	195
Sending Control Commands to SmartConnectors	195
Getting Connector Status	196
Sending Standard Flow-Control Commands	196
Tech Support Commands	198
Mapping Commands for Additional Data Fields	200
Managing SmartConnector Groups	203
Importing and Exporting SmartConnector Configurations	205
Importing a SmartConnector Configuration	205
Exporting a SmartConnector Configuration	206
SmartConnector Filters	206
Using Additional Data Fields	206
Upgrading SmartConnectors	209
Overview of the Upgrade Process	210
SmartConnector Upgrade Procedure	211
Rolling back to a Previous Version	212
Troubleshooting	212
Getting Status and Versions on Installed SmartConnectors	213
Consuming Events from Transformation Hub	213
 Chapter 7: Managing Notifications	 216
Managing Received Notifications	216
Managing Notification Groups	217
Managing Notification Destinations	219
Changing Notification and Acknowledgment Settings	220
Testing Notification Groups and Destinations	222
Managing Escalation Levels	222
 Chapter 8: Monitoring Events	 224
Monitoring Active Channels	224
Creating or Editing an Active Channel	224
Viewing Active Channels	228
Monitoring Events in the Active Channel	229
Full Search and Event Search on ArcSight Command Center	229
Using Views	231

Investigating Views	233
Viewing an Exploited Vulnerability	233
Viewing a Targeted Asset	233
Filtering an Active Channel	234
Filtering Active Channels with Inline Filters	234
Applying a Field Set to an Active Channel	236
Using an Active Channel Header	237
Sorting Events in the Active Channel	238
Adding, Replacing, or Removing a Column	239
Sizing, Showing, or Hiding Column Elements	241
Using Active Channel Menu Commands	241
Exporting Events to a File	243
Defining Grid Fields Options	245
Saving Copies of Active Channels and Filters	245
Best Practices to Optimize Channel Performance	246
Active Channels or Reports?	246
Active Channels or Query Viewers?	246
Active Channel Query Time Ranges	246
Active Channel Filters	247
Filtering on Indexed Fields	247
Filtering on Join Fields	247
Continuously Updating Time Parameters	247
Sorting by End Time or Manager Receipt Time	247
Sorting in Active Channels	248
Use of the “Live” Channel from Standard Content	249
Case Sensitive or Case-Insensitive Conditions?	249
I/O Subsystem Performance	249
Diagnostics: Start with Basic Channel Characteristics	249
Customizing Columns	250
Creating a Custom Column	250
Showing a Custom Column	251
Advanced Example: Creating a Custom Column with Velocity Template	251
Using Dashboards	252
Monitoring Dashboards	252
Creating or Editing a Dashboard	255
Adding a Data Monitor to a Dashboard	257
Adding a Query Viewer to a Dashboard	258
Dashboard Display Formats	259
Managing Dashboard Groups	260

Using Custom View Dashboards	261
Displaying Custom View Dashboards	261
Reverting to the Regular Dashboard View	262
Working with Custom View Dashboards	263
Arranging Custom View Dashboards	263
Loading a Background Image	264
Selecting a Previously Uploaded Background Image	264
Verifying the Background Image	265
Removing a Background Image	265
Custom View Dashboard Context Menu Options	265
Using Data Monitors	266
Creating a Data Monitor	266
Editing a Data Monitor	269
Deleting a Data Monitor	270
Managing Drilldowns from Data Monitors	270
Adding a Drilldown	270
Editing a Drilldown	275
Changing the Default Drilldown	275
Sorting or Changing the Order of Drilldowns	276
Removing a Drilldown	277
Moving or Copying a Data Monitor	277
Enabling or Disabling a Data Monitor	278
Overriding a Data Monitor's Last State	279
Managing Data Monitor Groups	279
Optimizing the Evaluation of Event Filters for Data Monitors	281
Requirement	282
Automating the Optimization of Filter Conditions	282
Tracing the Optimization	283
Disabling the Optimization Feature	284
Using Charts	284
Charting an Active Channel's Contents	285
Charting a Data Monitor's Contents	285
Exploring the Events Behind a Chart	286
Using Query Viewers	287
Graphing Attacks	287
Creating Static Event Graphs	287
Creating Live Event Graphs	288
Using a Data Monitor	288

Using an Active List	289
Event Graph Notes	289
Chapter 9: Selecting and Investigating Events in Active Channels	291
Selecting Events in the Active Channel	291
Showing Event Details and Rule Chains	291
Running Recon Searches	293
Investigating Session Events	294
Collaborating on Events (Event Annotation)	295
Annotating an Event	296
Mark Similar Events Fields	298
Annotation Preservation	299
Viewing Annotations for an Event	299
Creating or Editing Stages	300
Working with Event Payloads	301
Exporting Data Fields to a CSV File	303
Exporting Events to ServiceNow®	304
Getting Knowledge Base Articles	305
Chapter 10: Filtering Events	306
Creating or Editing a Filter	306
Creating and Editing an Inline Filter	307
Applying Filters	308
Moving or Copying Filters	310
Deleting Filters	310
Debugging Filters to Match Events	311
Importing and Exporting filters	312
Managing Filter Groups	313
Investigating Views	314
Using an Event Attribute to Show a New Filtered View	315
Refining a Filter with an Event Attribute	315
Filtering Out ArcSight Events	316
Adding an Event Attribute to a Filtering Condition	316

Modifying Views	317
Chapter 11: Queries	319
How Queries Work	319
Using Queries and Trends Together for Reports	319
Using Queries in Query Viewers	320
Building a Query	320
Query Settings	321
General Query Attributes	321
Query Fields	324
SELECT Query Fields	325
Query Structure (SELECT)	326
Applying Functions to SELECT Columns	327
GROUP BY Query Fields	328
Query Structure (GROUP BY)	329
Applying Time-Based Functions to GROUP BY Columns	330
ORDER BY Query Fields	331
Query Structure (ORDER BY)	332
Applying a Column Function to Order By	333
Sort Order	333
Query Conditions	333
Creating Conditions on a Field	334
Creating a Group Condition	335
Tips on Creating Conditions	335
Query Variables	335
Editing a Query	336
Example: Creating Asset-Related Conditions for Queries on Lists	337
Chapter 12: Query Viewers	340
Pre-Built and Custom Query Viewers	340
Standard Content	340
Custom Query Viewers	340
Customizing Query Viewers as Needed	341
inActiveList Conditions for Queries	341
Managing Query Viewers	341
Query Viewer Settings	343

Query Viewer Attributes	343
Query Viewer Fields	345
Sort Options	347
Baselines	347
Query Viewer Variables	348
Deleting a Query Viewer	349
Defining and Using Baselines	349
Why Baselines are Useful	350
Planning for Baseline Comparisons	351
Adding a Baseline	351
Comparing Displayed Results to a Baseline	352
Show or Hide Baseline Columns	353
Sort Baseline Data	354
Filter Baseline Data	354
Removing a Baseline	355
Managing Drilldowns from Query Viewers	356
Adding a Drilldown	356
Editing a Drilldown	360
Changing the Default Drilldown	360
Sorting or Changing the Order of Drilldowns	360
Removing a Drilldown	361
Viewing Query Viewer Results	362
Filtering Query Viewer Results	365
Viewing an Event or Resource Directly from the Query Viewer	367
Working with Query Viewer Results	367
Results in Table Format	368
"Analyze in Channel" Options on the Table View	368
Column Sort, Display, and Edit Options	370
Results in Chart Formats	371
Troubleshooting Query Viewers	372
Adding Query Viewers to Dashboards	372
Adding Query Viewers as Startup Views	373
Generating Reports from Query Viewers	374
Example Queries for Common Scenarios	375
Basic Analysis High Level Summaries	376
Analyst's First View of Events	376

How the Events Query Viewer is Built	377
Analyst's First View of Events	379
How the Events Query Viewer is Built	380
How the Events Query Viewer is Built	382
Drilldown Example	384
How the Console Builds Drilldowns	386
Non-Event Analysis Example	386
Baseline Analysis for Data Comparison	386
History Analysis Example	387
Chapter 13: Building Reports	388
Understanding the Reporting Workflow	388
Step 1 - Build a Query	389
Step 2 - Build a Trend Based on a Query	390
Step 3 - Build a Query Based on a Trend	390
Step 4 - Select or Design a Report Template	390
Step 5 - Create a Report	391
Step 6 - Run a Report	391
Step 7 - Archive and Maintain Reports	392
Managing Dependencies for Reports Resources	392
Creating or Editing a Report	393
Defining Report Attributes	394
Report Templates	395
Report Template Selection	395
Text Attributes	396
Preview	398
Binding Data to the Report	398
Binding Data to Charts	399
Selecting Data for the X-Axis on a Chart	400
Selecting Data for the X-Axis on a Chart	401
Selecting Data for the Y-Axis on a Chart	403
Selecting Data for the Z-Axis on a Chart (Optional)	405
Effect of Sorting on Bar Charts with Series Data	405
Specifying Top/Bottom Filters Aggregation Filters for a Chart (Optional)	407
Setting Display Options and Scale Formats for Charts	409
Binding Data to Tables	410
Specifying Fields for a Table	411
Enabling the Aggregation Tab for a Table	414

Setting Top/Bottom Counts in Table Aggregation Tab (Optional)	415
Setting Default and Custom Report Parameters	415
Adding Custom Parameters for Report Data	419
Displaying a Custom Parameter Prompt at Report Runtime	420
Adding or Removing a Prompt for Custom Parameters in the Report	420
Defining the Prompt in the Query's Condition Tab	423
Generating Reports with Central European, Cyrillic, or Asian Fonts	424
Creating Focused Reports	425
Using Report Templates	426
Applying a Standard Template to an Existing Report	427
Creating a New Report Based on a Template	428
Copying a Standard Template	429
Managing Report Template Groups	429
Editing a Template	430
End-to-End Reporting Examples	430
Example of Creating a Simple Report with the Wizard	431
Advanced Reporting Example Overview	433
Step 1 - Build the VPN Logins Outcome Query	433
Query Name and Other General Attributes	433
Fields to Include in Query Result	434
Query Conditions	436
Step 2 - Build the VPN Logins Outcome Hourly Trend	436
Step 3 - Filter the Trend Data (Login Attempts, Successes, Failures)	438
Step 4 - Create the VPN Logins Outcome Report on Trend Data	440
Choose a Template and Bind it to Result Data	440
Use Custom Parameters	442
Step 5 - Run the Report	444
Chapter 14: Running and Managing Reports	446
Running a Report	446
Running a Delta Report	450
Running Reports from a Grid View	451
Running a Rule Context Report	451
Running an Event Context Report	452
Running a Channel Report	452
Running a Query Viewer Report	453
Running Large or Complex Reports	456

Moving and Copying Reports	457
Managing Report Groups	458
Archiving and Scheduling Reports	459
Archiving a Report	459
Displaying an Archived Report	461
Scheduling Report Tasks	461
Scheduling Individual-Report Archiving	462
Scheduling Report Archiving by Resource Group	463
Standard Time Transitions	464
Viewing an Archived Report	464
Editing a Report Archiving Schedule	465
Editing Report Archiving Parameters	465
Deleting a Report Archiving Schedule	466
Chapter 15: Building Trends	467
How Trends Work	467
Snapshot Trend	468
Interval Trend	468
Query-Trend Relationships in Reporting	468
Managing Trends	470
Creating or Editing a Trend	470
Defining Trend Settings	471
Trend Attributes	471
Trend Schedule	475
Trend Parameters	476
Trend Actions (Add to Active List)	477
How Trend Actions are Useful (Summary Views and Rules)	477
Plan and Define Active Lists with Fields Mapped to Trend	478
Working with Trend Actions	478
Example: Populating Active Lists with Trend Results	480
Notes on Trend Action Behavior	482
Testing a Trend	482
Viewing Trend Data	483
Refreshing Trend Data	484
Disabling or Enabling a Trend	485

Deleting a Trend	485
Chapter 16: List Authoring	486
Required Settings for Large Lists	486
Creating or Editing an Active List	486
Viewing and Editing Active List Entries	492
Using Rules to Populate an Active List	494
Example Active List	494
Example Rule to Populate the Active List	495
Adding Events from a Channel to an Active List	497
Moving or Copying an Active List	497
Importing and Exporting an Active List	498
Deleting an Active List	498
Managing Active List Groups	499
Managing Session Lists	500
Creating or Editing a Session List	500
Editing Session List Entries	503
Moving, Copying, or Deleting a Session List	505
Exporting a Session List	505
Field Naming Restrictions	506
Chapter 17: Rules Authoring	507
Designing Rules	507
Rule Types	508
Managing Rules	509
Creating or Editing Rules	509
Moving or Copying Rules	510
Enabling and Disabling Rules	511
Viewing Rules and Their Correlation Events	512
Deleting Rules	513
Specifying Rule Conditions	513
Creating Rule Conditions	513
Adding Filter Conditions	515
Adding Asset Conditions	516
Adding Vulnerability Conditions	516

Adding Active List (InActiveList) Conditions	517
Creating Matching or Join Conditions	519
Editing or Deleting Join Data Field Conditions	521
Negating Event Conditions	522
Optimizing the Evaluation of Event Conditions	524
Automating Condition Optimization	525
Disabling the Optimization Feature	526
Tracing the Optimization	526
Specifying Rule Thresholds and Aggregation	528
Setting or Changing Rule Thresholds	528
Examples of Grouping Unique or Identical Field Values	529
Aggregation Time Criteria	530
Deleting Aggregation from a Rule	532
Managing Rule Actions	532
Adding, Editing, or Removing a Rule Action	533
Activating or De-activating a Rule Trigger	534
Enabling or Disabling a Rule Action	535
Threshold Triggering Options	535
Rule Actions Best Practices	538
Rule Actions Reference	539
Applying Rule Actions on Cases	546
Using a Rule to Create a Case	546
Using a Rule to Add to an Existing Case	547
Using a Rule to Create a Ticket in ServiceNow® ITSM	549
Converting Rule Types	550
Testing Rules	551
Debugging Rules	552
Verifying Rules with Events	553
Deploying Real-time Rules	555
Deploying a Rule	555
Removing or Un-deploying a Rule	556
Managing Rule Groups	556
Importing and Exporting Rules	558
Scheduling Rules	558
Scheduling a Rule Group	559
Scenarios for Using Scheduled Rules	560

Example of a Scheduled Rule (Badge Swipes and Logins)	561
Chapter 18: Identity Correlation	565
Understanding Session Correlation	565
Creating a Session List Rule	566
Using the Session List Output	568
Creating a Variable to Get Session List Data	568
Example: Using Session Lists to Correlate Session Data on User Logins	569
Step 1 - Create a Session List to Store Windows Sessions	570
Step 2 - Create Rules to Populate the Session List with Windows Logins	571
Rule 1: Triggers on Windows Session Logins	572
Attributes	572
Conditions	572
Aggregation	573
Actions	574
Rule 2: Triggers on Termination of Windows Sessions	575
Step 3 - Verify Rules	576
Step 4 - Use the Session List in a Report	578
Example: Using Active Lists to Correlate Users	580
Example Overview	581
Step 1 - Build and Populate the Active List with User IDs	582
Populating an Active List with User Data	583
Step 2 - Create a Rule that Uses Active List Values to Correlate User IDs	584
Attributes	584
Variable	585
Conditions	586
Aggregation	588
Actions	588
Step 2 - Create a Rule that Uses Active List Values to Correlate User IDs	589
Attributes	590
Variable	590
Conditions	591
Aggregation	593
Actions	593
Step 2 - Create a Rule that Uses Active List Values to Correlate User IDs	594
Attributes	594
Variable	595

Conditions	596
Aggregation	598
Actions	598
Chapter 19: Field Sets	600
Creating a Field Set	601
Field Set Editor: Attributes Tab	602
Field Set Editor: Fields Tab	602
Using the Fields & Global Variables Subtab	603
Using the Field Sets Subtab	604
Using the Local Variables Subtab	604
Field Set Editor: Local Variables Tab	605
Adding Custom Columns to the Field Set	605
Renaming a Column Using an Alias	606
Editing a Field Set	607
Sharing a Field Set	607
Deleting a Field Set	608
Resources That Use Field Sets	608
Chapter 20: Global Variables	609
Remote Variables Processing	609
Global Variable Dependencies	609
Navigating to Global Variables	610
Creating or Editing a Global Variable	610
Global Variable Editor: Attributes Tab	612
Global Variable Editor: Parameters Tab	612
Global Variable Editor: Local Variables Tab	613
Moving, Linking, or Deleting Global Variables	614
Promoting a Local Variable to a Global Variable	614
Adding a Global Variable to a Resource	617
Accessing a Global Variable Using the CCE	617
Adding Global Variables to an Active Channel	619
Adding a Global Variable to a Data Monitor	620
Adding a Global Variable to a Field Set	621
Chaining a Global Variable	622

Chapter 21: Case Management and Queries	624
Creating or Editing a Case	625
Locking and Unlocking Cases	627
Entering Case Attributes	628
Entering Case Descriptions	631
Entering the Case Security Classifications	632
Entering Follow Up Items for the Case	633
Entering Attack Mechanism Information	633
Entering Attack Agent Information	634
Entering Incident Information	635
Entering Vulnerability Information	636
Entering Miscellaneous Information	636
Using the Case's History Panel	637
Working with Events in Cases	640
Creating or Updating a Case from Displayed Events	641
Using the Case Events Panel	642
Viewing a Case's Events in a Channel	643
Including Base Events Through a Rule	643
Copying Event Details from Case to Case	644
Deleting Events from a Case	644
Attaching a File to a Case	644
Attaching a Data Monitor, Dashboard, or Query Viewer to a Case	647
Viewing a Case Attachment	648
Editing a Case Attachment	648
Best Practices on Attaching Files to a Case	649
Closing a Case	650
Deleting a Case	650
Granting Permission to Delete Cases	651
Moving or Copying a Case to a Group	652
Finding Cases	652
Viewing a Case's Internal Audit Events	652
Managing Case Groups	653
Viewing Group Cases in a Grid View	655
Running Case Queries	656
Creating a Report from a Case	656

Running Case Reports and Setting Default Parameters	657
Customizing the Case Report	659
Customize Selected Case Query	660
Customize Selected Case Report	661
Add a Server Property for the New Report URI	662
Using External Case Management Systems	662
Exporting Cases to ServiceNow®	663
 Chapter 22: Integration Commands	 664
What are Integration Commands?	664
Supported Command Types	664
Local Scripts and Commands to Other Applications	665
How Integration Commands Work	665
Planning Checklist and Workflow	666
Navigating to Integration Command Resources	667
Defining Commands	668
Script Commands	669
URL Commands	671
Connector Commands	672
Adding and Editing Command Parameters	673
Removing a Command Parameter	675
Using Configurations to Group Commands	675
Configurations Attributes	677
Configurations Contexts	678
Configurations Commands	679
Configuration Targets	680
Adding a Target to a Configuration	680
Editing Targets in a Configuration	680
Removing Commands from a Configuration	681
Specifying Targets	681
Target Attribute	681
Target Integration Parameters	682
Authorization and Authentication Settings	682
Setting User Login Parameters	683
Setting Login Credentials	683
Setting Login Credentials on Target Servers	684

Setting Logins and Other Parameters to Prompt for Values at Runtime	684
Running Integration Commands	685
Entering/Saving Command Parameters at Runtime	685
Using the Recon Integration Commands	686
ArcSight Logger Search Commands	688
Logger Integration Commands	688
Enabling Integrated Logger Searches	689
1. Set Up Logger Command Targets	690
2. Set Up the Logger Command Configuration	690
3. Set Up Users for Logger Access	691
Example of Running a Logger Quick Search	692
Network Tools as Integration Commands	692
More Integration Examples	694
 Chapter 23: Knowledge Base Authoring	 698
Managing Knowledge Base Articles	698
Managing Knowledge Base Article Groups	701
Associating Knowledge Base Articles	702
 Chapter 24: Finding Resources	 704
How Fields are Indexed	704
Using Text Search Syntax	705
Using the Search Field on the Console Tool Bar	708
Using the Search Result Columns	710
Locating Resources on the Navigator Tree	710
 Chapter 25: Managing Resources	 711
Working with Resource Groups	711
Adding or Editing a Resource Group	711
Using the Categories Tab for Asset Groups	712
Moving, Copying, Linking, and Deleting Resources	712
Locking and Unlocking Resources	713
Selecting Resources	714

Visualizing Resources	714
Graphing Resources	715
Using Graphs	716
Configuring Resource Graphs	717
Viewing Resources in Grids	717
Validating Resources	718
About Valid and Invalid Resources	718
Fixing and Validating Resources	718
Troubleshooting Requirements for Valid Resources	721
Resource Validation During Upgrade or Package Import	723
Extending Audit Event Logging	724
Common Resource Attribute Fields	725
Common	725
Assign	726
Saving Copies of Read-Only Resources	726
Managing File Resources	727
Uploading Files and Creating a File Resource	727
Working with Files	728
 Chapter 26: Managing Packages	 730
Creating or Editing Packages	731
About Locked Packages	735
Adding Resources from the Resource Navigator	735
Supported Packages for Content Synchronization	735
Exporting Packages	737
Importing Packages	737
Best Practices for Importing Packages	738
Importing Packages Created by Other Users	739
Backing Up and Restoring with Packages	740
ID Checking During Import	740
Package Modifications	741
List Data	741
Backup and Restore Summary	742
Installing or Uninstalling Packages	742
Deleting Packages	744

Removing Resources from Packages	745
Resolving Package Conflicts	745
Chapter 27: Using Threat Detector	747
Threat Detector Overview	747
What Pattern Detection Provides	747
Pattern Components	748
How Threat Detector Works	749
Threat Detector Life Cycle	750
Creating or Editing a Profile	750
Specifying Actions	753
Creating Local Variables	756
Adding Notes	757
Deleting a Profile	758
Taking a Snapshot	758
Analyzing Snapshots	760
Exploring a Snapshot	760
Arranging Elements in Graphic View	762
Scheduling a Snapshot	763
Re-opening a Snapshot	764
Deleting a Snapshot	765
Investigating Patterns	765
Investigating Patterns in the Snapshots View	765
Investigating Patterns in the Patterns View	767
Viewing Patterns with Filter	768
Inspecting Patterns	769
Creating Rules from Patterns	770
Annotating Patterns	772
Deleting a Pattern	773
Threat Detector Usage Guidelines	773
Establishing a Baseline of Normal Patterns	773
Using Threat Detector in Routine Operations	773
Performance Considerations	774
Adjusting Threat Detector Memory	774
Chapter 28: Actors	775

Configuring Actors	775
Permissions Required to Use Actor-Related Data	777
Viewing Actors on the	778
Viewing an Actor in the Actor Editor	779
Viewing Actor Account Attributes	780
Viewing Actor Role Attributes	781
Viewing Actors in an Actor Channel	781
Sorting Fields in Actor Channels	782
Actor Channel Options	783
Filtering Actor Channels	783
Adding a Local Filter to the Actor Channel Resource	784
Creating an Inline Filter	785
Managing Actor Channels	785
Investigating Actors	786
Running Context Reports from an Actor Channel	786
Investigating an Actor from an Event Channel	788
Actor Context Reports in Standard Content	789
Creating and Editing Actors for Testing Purposes	790
Important Points to Consider About Making Manual Changes to Actors	790
Creating Actors for Testing Purposes	790
Editing Actors for Testing Purposes	792
Deleting Actors	793
Leveraging Actor Data Using Variables	793
Creating an Actor Global Variable	794
Creating an Actor-Based Variable in Another Resource	794
Creating and Using Category Models	795
Memory Recommendations for Using Category Models	795
Creating Category Models	796
Creating Actor-to-Actor Category Models	797
Creating Actor Attribute Category Models	800
Creating User-Defined Category Models	802
Managing Category Models	804
Viewing Category Models in Graphs	805
Leveraging Category Model Data Using Variables	808
Chapter 29: Reference Guide	810

Access Control Lists	810
Resource ACLs	810
Active Channels	812
Active Channel Views	812
Active Channel Headers	813
Comparisons	814
Active Channel Views for Assets and Cases	814
About Actors	814
How the Actors Feature Works	817
About the Actor Model Import Connector	818
Troubleshooting Errors with Actor Model Imports	819
Active Lists	820
Uses of Active Lists	821
Active Lists for Long-Term State Retention	821
Optimize Data with Hash-Based Active Lists	821
Active List Monitor Events	822
Active Lists with Values	823
Using Variables to Retrieve Data from Active Lists with Values	823
Example: Active List with Values to Store Directory Information	824
Create an Active List	824
Populate the Active List	824
Correlate Information Stored in UserRoles List	825
Administrator	827
Advanced Editor	828
Aggregation	829
ArcSight Console	829
Assets	830
Assets Tab	830
Zones Tab	831
Networks Tab	831
Categories Tab	831
Vulnerabilities Tab	832
Locations Tab	832
Attack	833
Audit Events	833
Audit Events Common to Most Resources	834

Active Channel	835
Active List	835
Actor	835
Archive	836
Authentication	836
Authorization	837
Backpressure Audit Events	837
Connector Connection	838
Connector Exceptions	838
Connector Login	840
Connector Registration and Configuration	840
Content Management	841
Created Ticket on External System	842
Dashboard	842
Data Monitors	842
Distributed Correlation	845
Aggregator Audit Events	845
Correlator Audit Events	846
DCache (Distributed Cache) Audit Events	846
MBus (Message Bus) Audit Events	846
Persistor Audit Events	847
Distributed Event Forwarding	847
Global Variables	847
Group Management	848
License Audit	848
Logger Component	849
Alerts	850
Archives	853
Certificates	854
Peers	855
Saved Searches	856
Searches	857
Search Filters	858
Storage Groups	858
Storage Volume	859
Manager Activation	859
Manager External Event Flow Interruption	859
Mark Similar	859
Status Monitor Events	860

Active Channel Statistics	860
Active List Statistics	860
Asset Statistics	861
Data Monitor Statistics	862
Transformation Hub Statistics	862
Filter Engine Statistics	863
Main Flow Statistics	863
Notification Statistics	863
Threat Detector Statistics	864
Report Statistics	864
Resource Framework Statistics	864
Rules Engine Statistics	865
Session List Statistics	866
Session Management Statistics	866
SmartConnector Flow Statistics	866
Notification	867
Notification Acknowledgement, Escalation, and Resolution	868
Notification Testing	868
Threat Detector	868
Query Viewers	869
Reports	869
Repository Audit Events	869
Resource Quota	870
Rule Actions	870
Rule Activations	871
Rule Firings	872
Rule Warnings	872
Rules Scheduled	872
Scheduler Execution	873
Scheduler Scheduling Tasks	873
Scheduler Skip	873
Session Lists	874
Transformation Hub	874
Trends	875
Trend Partitions	875
User Login	876
User Management	876
Base Queries	876

Batching	877
Cases	877
Case Groups	877
Categories	878
Object Category	879
Behavior Category	881
Outcome Category	882
Device Group Category	883
Significance Category	883
Technique Category	884
Asset Categories	886
Event Categories	886
Collaboration	886
Common Conditions Editor (CCE)	887
Editor Features	888
Condition Tree Command Buttons	890
Condition Tree Context Menu Commands	892
Adding Conditions	895
Search Box to Find Fields in the List	897
Field Comparisons with Variable or Static Values	898
Using Field Sets	899
Adding or Removing Global Variables Using the CCE	901
Testing for Zone Relevance	902
Conditional Statements	903
Conditions	904
Parameterized Conditions	904
Content	905
Content Packages	906
Custom Content	906
SmartConnector Content	906
CORR-Engine	907
Correlation	907
Correlation Formula	907
Correlation Rule	908
Customers	909

Dashboards	909
Dashboard Context Menu Commands	910
Data Fields	910
Attacker Group	911
Connector Group	915
Category Group	919
Destination Group	920
Device Group	924
Device Custom Group	929
Event Group	932
Event Annotation Group	939
File Group	942
Final Device Group	942
Flex Group	945
Geographical Attributes	945
Manager Group	946
Old File Group	946
Original Connector Group	947
Request Group	951
Source Group	952
Target Group	955
Threat Group	958
Resource Attributes	959
Data Monitors	959
Asset Category Count Data Monitor	960
Event Correlation Data Monitor	961
Event Graph Data Monitor	964
Geographic Event Graph Data Monitor	965
Hierarchy Map Data Monitor	966
Hierarchy Map Features	967
Use Cases	967
Defining a Hierarchy Map Data Monitor	968
Adding Variables	969
Specifying the Source Node Identifiers	970
Hierarchy Levels and Group Delimiters	970
Specifying Group Attributes	971
Hierarchy Map Display and Visualization Controls	972
Map Display and An Example	972

Labels, Size, and Color Controls	973
Selecting Colors for the Blocks	975
Hourly Counts Data Monitor	976
Last N Events Data Monitor	977
Last State Data Monitor	978
Last State Data Monitor Parameters	979
Options for Table and Tile Views	980
Table View (Color Chooser and Remove Entry)	980
Tile View (Customize View)	980
Moving Average Data Monitor	982
Rules Partial Match Data Monitor	985
Statistics Data Monitor	986
System Monitor Data Monitor	988
System Monitor Attribute Data Monitor	989
Top Value Counts Data Monitor	990
Troubleshooting	992
Data Monitor Expressions	993
Supported Data Monitor Expression Operators	993
Supported Data Monitor Expression Functions	993
Device	994
Event Inspector	994
Events	995
Event Annotation Fields	997
Event Categorization	997
Event Handling Stages	998
Field Sets	999
Filters	1000
Filtering Options	1000
Global Variables	1002
Grid View	1002
IP Address Ranges	1002
Inspect/Edit Panel	1003
Job Scheduler	1003
Viewing all scheduled jobs	1004
Troubleshooting Tips	1005

Knowledge Base	1005
Logical Operators	1006
Managed Security Service Providers (MSSPs)	1009
Manager	1009
Navigator Panel	1010
Notifications	1010
Notification Operation	1010
Testing Notification Escalations	1011
Notification Destinations	1012
Notification Acknowledgements	1012
Packages	1012
Threat Detector	1013
Pattern Concepts	1014
Discovering Patterns	1014
Pattern Analysis	1015
Initial Phase	1015
Routine Pattern Processing	1015
Workflow Management	1015
Pattern Analysis	1016
Pattern Disposition	1016
Threat Detector Expertise	1016
Workflow	1017
Visualization	1017
Applications	1017
Payload	1018
Prioritization Fields	1018
Priority Calculations and Ratings	1019
Priority Elements	1022
Priority Operators	1023
Priority Rating	1024
Queries	1025
Queries and Trends	1025
Building and Running Queries	1025
Query Viewers	1026
Reference Pages	1027

Reports	1028
Working with Report Templates, Queries, and Trends	1028
Viewing and Managing Reports	1029
Archived Reports	1029
Report Groups	1029
Delta Reports	1030
Report Parameters	1030
ArcSight-Provided Reports	1031
Report Templates	1032
Resources	1032
Resource Attributes	1032
Rules	1034
Loading Rules	1035
Automatically Disabled Rules	1035
Rules Processing and Correlation	1037
Rule Groups	1039
Scheduled Rules	1040
Rule-triggering Timing	1040
Rule Chains	1041
Variables	1041
Rule Actions	1041
Rule Conditions	1041
Rules Editor	1043
Saved Searches	1043
Schema	1043
Avoiding Field Naming Collisions	1044
Event Fields	1045
Precise Event Categorization	1045
Search Filters	1046
Send Logs	1047
Guidelines for Using the Send Logs Utility	1047
Options for Running Diagnostics and Sending Logs	1048
Starting the Send Logs Wizard on the ArcSight Console	1048
Session Correlation	1049
Why Session Correlation Matters	1049

Session Lists	1049
SmartConnectors	1050
Operational Status	1051
Configuration	1051
Zones	1052
Upgrading	1053
Filtering	1053
SMTP	1053
Sortable Field Sets	1054
Sorting Columns in Grid Views	1055
Threat	1055
Threat Evaluation	1056
Evaluation Process	1056
Evaluation Definitions	1056
Maintaining Model Confidence	1057
Using Threat Evaluation Information	1057
Limitations and Workarounds	1058
Thresholds	1058
Time Error Correction	1059
Timestamps	1059
Timestamps for Security Events	1059
Timestamps for Resources	1060
Timestamp Variables	1060
Inclusive Timestamps	1061
Time Zone Correction	1061
Understanding Trends and Queries	1061
User Groups	1062
Users	1063
User Types	1063
Variables	1064
About Remote Variables	1065
About Functions	1065
Local and Global Variables	1066
Variable Definition Fields	1066
Alias Functions	1067

Arithmetic Functions	1067
Category Model Function	1070
Condition Functions	1072
Group Functions	1073
IP Address Functions	1074
List Functions	1075
String Functions	1076
Timestamp Functions	1077
Type Conversion Functions	1080
Value List Functions	1084
Using Functions: Examples with Lists	1085
Getting Login Session Data from a Session List	1085
Extracting a List Element from an Active List	1086
Variable Availability and Contexts	1087
Variable Functions for In-Memory Operations	1087
Velocity Templates	1088
Velocity Application Points	1088
Using Velocity Expressions to Retrieve Values from Event Fields or Variables	1089
Retrieving Values from Event Fields	1090
Using Variables in a Velocity Expression	1090
Using Velocity Expressions in Rule Actions	1090
Example of Rule Action that Uses Velocity Expressions to Retrieve Values	1090
Velocity References for Reports	1091
More Velocity Template Examples	1094
Velocity Template Usage Tips	1095
Views	1096
View Types	1096
Dashboards	1097
Other Views	1098
Vulnerabilities	1098
Vulnerability Groups	1098
Standardized Vulnerability Tracking	1099
Web Browsers	1099
Browser Preferences for HTML Displays	1099
Browser Preference Overrides for Specific Features	1100
Publication Status	1101

Send Documentation Feedback	1102

About this PDF Version of Online Help

This document is a PDF version of the online help. This PDF file is provided so you can easily print multiple topics from the help information or read the online help in PDF format. Because this content was originally created to be viewed as online help in a web browser, some topics may not be formatted properly. Some interactive topics may not be present in this PDF version. Those topics can be successfully printed from within the online help.

Chapter 1: Getting Started

Welcome to ESM and the ArcSight Console.

ESM is a comprehensive software solution that combines traditional security event monitoring with network intelligence, context correlation, anomaly detection, historical analysis tools, and automated remediation. It consolidates and normalizes data from disparate devices across your enterprise network in a centralized view.

Starting the ArcSight Console

Start the as you would any other application.

Start the Console:

Depending on the chosen shortcuts during installation, start the using any of these methods:

- Using the desktop icon
- Selecting from the system tray
- Selecting from the Start menu

Alternatively, open a command window in the 's bin directory and type

```
arcsight console
```

Log in:

The login mechanism varies according to the type of authentication you have set up during installation.

- If you are using SSL authentication, set it up and import the certificate as described in the *ESM Administrator's Guide's* "SSL Authentication" section, the topic "Setting Up SSL Client-Side Authentication on ArcSight Console." After the certificate is imported, you can start the without entering a user ID or password.
- If you are using password authentication, log in with your user ID and password. Certificates are imported automatically. See the *ESM Administrator's Guide's* "Configuration" section, in the topic "Managing Password Configuration" for more information.
- If you have selected "Password or SSL Authentication," you choose which way to log in, each time.
- If ESM is configured to use OSP Client Only Authentication or External SAML2 Client Only Authentication, the **OSP Client Login** button is automatically enabled. You will be redirected to your One SSO Provider (OSP) to authenticate to the console.



Note: If either of these authentication methods are in use, Open Text recommends running only one console instance per workstation at a time. Otherwise, authentication will not work as expected. For example:

- If a console instance is running and you have authenticated to the console, if you open a second console instance the same user that is logged in to the first instance is logged in to the second instance.
 - If you open two instances of the console before you authenticate, two browser sessions are opened for OSP login and the login attempt will fail. You must close all consoles and browser windows before you attempt to log in again.
- If you are using OSP Client Only Authentication on the server, you must import the OSP certificate into the ArcSight Console keystore. See the “One SSO Provider (OSP) Authentication” topic in the *ESM Administrator’s Guide*.
 - If you are using FIPS and using a browser, make sure that browser is configured for FIPS. See the “Configure Your Browser for FIPS” topic in the *ESM Administrator’s Guide*.

Quick Start Tools and Standard Content

The [ArcSight Console](#) serves as the control point for administrators to:

- Configure ESM content and resources
- Manage, monitor, and respond to network security issues across the enterprise

A Network Model Wizard is provided to facilitate the process of describing network devices and assets in ESM. For more about the Network Model wizard and instructions on how to use it, see ["Populating the Network Model Using the Wizard" on page 132](#).

A set of coordinated *resources* (filters, rules, dashboards, reports, and so on) is provided to address common security and management tasks. The set of standard content is designed to give you comprehensive correlation, monitoring, reporting, alerting, and case management out of the box, with minimal configuration required on the [ArcSight Console](#).

For information about standard ArcSight administration or system content, see the [ArcSight Administration and ArcSight System Standard Content Guide](#).

Use Cases

Use cases are special groupings of related ArcSight content that address specific security issues and business requirements.

Use cases provide an integrated -based alternative for viewing and interacting with resources to the standard one-resource-at-a-time viewing method offered in the Resource tree of the [ArcSight Console](#).

Navigator panel. You can configure shared resources in a single operation, and export related resources in an ArcSight Resource Bundle (arb) for use in other ArcSight instances.

OpenText provides use cases for some of the standard content that is installed with ESM and for additional content (Security Use Cases) provided through the Marketplace. Each Security Use Case comes with its own documentation that provides information about how to install, configure, and use the use case.

Use case configuration requires having a network model in place. Model your network first as part of the initial configuration of ESM. Follow instructions in ["Modeling the Network" on page 119](#).


Chapter 2: Working in the

In addition to the capabilities built into the , the Console itself is a tool with its own characteristics and specialized controls. The Help topics in this section describe the basics of using tools and controls to make the most of its features.

Navigating

Use the Navigator panel on the to locate and manage security resources, and the Viewer and Inspect/Edit panels to analyze resource data and view or adjust the attributes of the resources producing the data.

The Navigator panel showing the Dashboards resource tree



The resources available in the Navigator panel can be affected by permissions set for your user type.

On the Navigator panel, you can:

- Choose a group or a specific resource from the resource tree.
- Expand (+) and collapse (-) resource groups to locate particular subgroups or individual resources. You can also use the keyboard **right arrow** key to expand and **left arrow** key to collapse the Navigator resource trees.
- Right-click groups or individual resources to choose from their context menus.
- See a list of the last 10 resources you have recently viewed and add resources to your favorites list

Use the Viewer or Inspect/Edit panels to see or act on the results of the context menu commands.

Navigator Panel Resource Tree

Resource Tree on the Console's Navigator Panel

Tree	Icon	Resource
Active Channels		Create, modify, and delete security-event views that actively and continuously evaluate the events they display, on the basis of time and other filter conditions. This view also includes the Field Sets resource tree for managing named field sets. See "Monitoring Events" on page 224 .
Actors		Map humans or agents to activity in applications and on the network, and identify actors behind events. See "Actors" on page 775 .
Assets		Security-sensitive devices and device groups installed in your enterprise, and the known exposures to potential threats those devices may represent. Assets also includes the related network, zone, location, category, and vulnerability information you use to manage network devices. See "Modeling the Network" on page 119 .
Cases		Track enterprise security incident cases, by status and priority. See "Case Management and Queries" on page 624 .
Connectors		Manage the SmartConnectors installed at your enterprise. See "Managing SmartConnectors" on page 175 .
Customers		Manage resources that represent the security concerns of particular MSSP (Managed Security Services Provider) clients. See "Managing Customers" on page 172 .
Dashboards		Various event data monitors and their library of supporting resources. See "Using Dashboards" on page 252 .
Field Sets		Define subsets of available data fields so you can quickly focus a grid view, an Event Inspector, or other field arrays on a particular context. See "Field Sets" on page 600 .
Files		The Files resource tree, when populated, lists files saved as resources on the Manager. This makes them accessible to all users of the system who are authorized for such access. File resources include Case file attachments, templates, and general-purpose shared files. See "Managing File Resources" on page 727 .
Filters		Event filtering definitions, organized in groups. See "Filtering Events" on page 306 and "Managing Filter Groups" on page 313 .
Integration Commands		Application integration resources used to configure and launch commands, tools, and views in custom and third party applications and other ArcSight products from within the . Provides the ability to configure custom scripts, URLs, and Connector commands, and integrate them into the Console UI in various contexts. Leverages velocity expressions and the UI contexts for pulling the content of event data, for example, as command parameter values.
Knowledge Base		A database of articles and groups of articles that aid problem-solving, analysis, and operation. See "Getting Knowledge Base Articles" on page 305 and "Knowledge Base Authoring" on page 698 .

Resource Tree on the Console's Navigator Panel, continued

Tree	Icon	Resource
Lists		Active Lists are lists of active source and target IP addresses of interest, as defined by enterprise rules. See "List Authoring" on page 486 for more information. Session Lists are similar to active lists, but are optimized for time-based queries and monitoring of rule-driven combinations of event attributes or custom fields. See "Identity Correlation" on page 565 for more information.
Notifications		Destinations and settings for the automatic messages that alert you to pre-defined situations or events. See "Acknowledging Notifications" on page 77 and "Managing Notifications" on page 216 .
Threat Detector		Formerly known as Pattern Discovery. Profiles to capture, and snapshots of, potentially threatening event patterns. See "Using Threat Detector" on page 747 .
Query Viewers		A resource for defining and running SQL queries on other ESM resources (independent of reports), including trends, assets, cases, connectors, events, and so forth. Each query viewer contains an SQL query along with other logic for establishing and comparing baseline results, analyzing historical data to find patterns in network activity, and performing drill-down investigation on a particular aspect of the results. Query viewers can use the same queries as reports do, but can be run independently of them. See "Query Viewers" on page 340 .
Reports		Definitions for, and archived output from, various activity reports. See "Building Reports" on page 388 and "Running and Managing Reports" on page 446 .
Rules		Rules and groups of rules created for isolating, analyzing, and responding to events. See "Rules Authoring" on page 507 .
Saved Searches		Saved Searches are created on the ArcSight Command Center. Refer to the ArcSight Command Center Users Guide for information on how to create and save searches. This resource is displayed on the ArcSight Console for packaging and content synchronization purposes. See "Managing Resources" on page 711 and "Managing Packages" on page 730.
Search Filters		Search Filters are created and used on the ArcSightCommand Center. Refer to the ArcSight Command Center Users Guide for information on how to create searches, then save them as filters. This resource is displayed on the ArcSight Console for packaging and content synchronization purposes. See "Managing Resources" on page 711 and "Managing Packages" on page 730.
Stages		Workflow and annotation features for real-time analyst collaboration on security events.
Use Cases		Resource collections that address common security issues and business requirements. When use cases are installed, a Use Case tab is displayed in the Navigator panel. A wizard is available for configuration of the use case resources. Instructions for using the wizard are provided in the documentation provided with the specific Use Case.
Users		ArcSight users and user groups. See "Managing Users and Groups" on page 99 and "Managing Permissions" on page 107 .

Batch Editing

You can make common edits to multiple case or SmartConnector resources by selecting a set of either type in the Navigator panel and changing their common fields in the Case or Connector Editor.

Batch-Editing Cases or Connectors

Where:

- Navigator > Resources > Connectors, or
- Navigator > Resources > Cases

To batch-edit cases or connectors:

1. **Ctrl+click** or **Shift+click** to select a set of individual cases or SmartConnectors in their respective resource trees.
2. **Right-click** the selected items and choose **Edit**.
3. Make changes to the appropriate common fields, such as **Description** or **Owner**.
4. Click **Apply** to record your changes and leave the editor open, or click **OK** to save and close. Saving affects only the fields you have changed, in each of the selected resources.

Locking Case Groups

Use the **Lock Case** check box to lock and unlock cases in batches. See ["Viewing Group Cases in a Grid View" on page 655](#).



Note: If a rule action is configured to update a case, and the case is locked at the time the rule triggers, then the case will not be updated. See ["Applying Rule Actions on Cases" on page 546](#).

SmartConnector Reminders

Batch changes affect only default configurations, not alternates. However, you can add new alternate configurations by batch editing.

Note that if you make changes under the **Filters** tab, the entire tab's contents are saved to the selected SmartConnectors.

You can batch-edit connectors only of the same version.

Reconnecting to the Manager


If your loses its connection to the Manager, a dialog popup enables you to **Retry** the connection, **Relogin**, or to **Cancel** the connection. Try these options in this order.

A connection to the Manager cannot be re-established if the Manager is restarted or if a network problem prevents communication with the same Manager. In such cases, click **Cancel** and start the Console again, using an appropriate Manager host name.

Changing the Console Display

You can change the look and feel of the Console to better display information, focus on particular panels, or hide information not of interest. You can switch to a dark theme, resize the Console, float or dock Console panels, apply translucency to a floating panel, and show or hide the menu bars, tool bars, and various displays.

What do you want to do?	Here's how:
Switch from default to dark theme	From the View menu, select Themes. You have two options: • Default is the daylight theme appropriate for a lighted room. • Dark theme is appropriate for a dark room environment to reduce glare. If you switch the theme, log off, then log back in. Note: After you have used the dark theme for a while, you may notice that the labels on the tabs are no longer legible. If so, exit the Console and log back in.
Resize the	• To expand to the whole screen, click the Maximize icon at the top-right corner of the window. • To collapse the Console, click the Minimize button or drag the corners of the Console to resize it. • To resize any panels, drag and drop any panel dividers.
Show or hide menu bars and tools	Right-click the Menu bar area of the Console and use the context menu to enable (check) or disable (clear) each component.
Show or hide the status bar 	Click the Status Bar button on the toolbar, or on the Window menu, choose Status Bar .
Show or hide the Navigator panel 	Click the Navigator button on the toolbar, or on the Window menu, choose Navigator Panel .

What do you want to do?	Here's how:
Show or hide the Viewer panel 	Click the Viewer button on the toolbar, or on the Window menu, choose Viewer Panel .
Show or hide the Inspect/Edit panel 	Click the Inspector button on the toolbar, or on the Window menu, choose Inspect/Edit Panel .
Float a Console panel 	Click the Float/Dock button on the panel header, or right-click the panel header and choose Float Panel . You can apply translucency once a panel is floated.
Apply translucency to a floating Console panel	Float the panel first before applying translucency. Move the Translucency slider on the panel header. 
Dock a Console panel 	Click the Float/Dock button on the panel header, or right-click the panel header and choose Dock Panel .
Close a Console panel 	Click the Close button on the panel header, or right-click the panel header and choose Close Panel .

Changing User Preferences

You can change several Console characteristics to suit your security needs, working style, or personal preferences. You reach the Preferences dialog box through the **Edit>Preferences** menu command.

Changing Your Password

Administrators create users and assign passwords. After logging in with your administrator-created password, you must change it for security reasons.



Note: You can change your password only if your ArcSight installation is configured to use built-in password authentication. Contact your system administrator for instructions on how to change passwords on ArcSight systems that use RADIUS SecurID or SSL authentication.

Where: **Edit > Preferences > Password**

1. Enter your old password, new password, and confirm the new password.
2. Click **OK**.

By default, passwords require a minimum of 6 characters, can contain a maximum of 20 characters, and can contain numbers and/or letters. Ask your system administrator about any special requirements for your site. For information on password restrictions, see "Managing Password Configuration" in the [ESM Administrator's Guide](#).



Note: If you are an administrator, you can change other ESM users' passwords. See ["Resetting User Passwords" on page 104](#).

Setting Default Editors and Viewers

You can set the default editors and viewers to use for text, HTML, and packet payloads. For example, use the HTML editor when editing the Knowledge Base and the Web browser for reports.

Where: Edit > Preferences > Programs

Program Preferences

Program Preference	Value
Preferred Text/HTML Editor	Enter the complete path to your preferred text or HTML editor, or click the Browse button to locate the editor.
Preferred Web Browser	Enter the complete path to the preferred Web browser or click Browse to locate the executable. Use your preferred Web browser to display HTML files such as custom view dashboards, reports, knowledge base articles, and so on. For an updated list of supported products, refer to the <i>ESM Support Matrix</i> on the ESM documentation page. This matrix includes the supported Web browsers for all currently-supported ESM versions.
Preferred Payload Viewer	Enter the complete path to your preferred packet-payload viewer or click the Browse button to locate one.
Text to PCAP Converter	Enter the complete path to your preferred packet-payload PCAP converter or click the Browse button to locate one.

Changing Global Options

You can make the Inspect/Edit panel open as a docked window inside, or as a floating window outside, the Console. You can do the same with all child windows as a class.

Where: Edit > Preferences > Global Options

Refer to the following table for available settings:

Global Options for

Global Option	Description
Font	Set global preference for font face, size, and style used throughout the Console, except on windows or views where you can set fonts specific to those Console elements. (For example, you can set fonts specific to Grid views as detailed in the next topic.) Click into the Font field to get the drop-down menu arrow.  Click the arrow to bring up the Fonts dialog. Set the Font, Size and Style.
Launch editors in a floating window	Open all editors in a floating window. If deselected, all editors appear in the Inspect/Edit panel. If you select this option, you can still float or dock the windows.
Allow multiple editors of the same type	Permit more than one resource editor to be opened simultaneously for a given resource type (for example, opening three instances of the Filter Editor at once). Enabling this option is very useful for analysts and persons implementing security solutions, but might be inappropriate for operators or other persons who should have less-extensive editing access.
Allow multiple event inspectors	Display details of multiple events in their respective Event Inspector tabs on the Inspect/Edit panel. If de-selected (the default), you can only view event details one event at a time. For more information about the Event Inspector, see "Inspecting and Editing" on page 68 and "Event Inspector" on page 994.
Allow Bulk Delete	Delete multiple resources without any dependency warnings. If de-selected, you can still delete multiple resources but you see a warning if there are any resource dependencies.
Create independent floating windows	Independently float new windows that are children of another window such as the Viewer panel. This is the default. When enabled, you can choose a window's name from the list at the Window>Floating command, or toolbar button to bring it forward: 
Auto Relogin	Automatically log in again after logging out of the Console.


Global Options for , continued


Global Option	Description
Use system defaults for dashboard background	When this option is selected, your system defaults are used for all Dashboard backgrounds. You can customize views for dashboards for display on the Web browser. See "Using Custom View Dashboards" on page 261 .
Show print preview dialog	Display a preview of the printable page when you choose to print a resource definition, for example, a rule definition. This preference is selected by default. Print preview options include Print, view each printable page (as applicable), and zoom in or out of the previewed page. For more information about printing, see "Printing from the Console" on page 93 .
Set Help dialog size (Width,Height)	The Help display window defaults to width of 910 x length of 650 pixels. Set Help dialog size (Width,Height) 910,650 You can specify a different default Help window display size here. To do this, enter a new window size (for example: 750,900), then press the Enter key. Note: Press Enter after setting the new display size, and then <i>also</i> click Apply or OK to save all preference settings. If you do not press Enter, the new window size setting cannot be saved even if you click Apply or OK.

Setting Dialog Options

Purpose: Part of Global Options, Dialog Options is where you define the behavior of dialog boxes for system messages. System messages are classified into error and informational or warning messages.

Where: Edit > Preferences > Global Options





Refer to the following table for available settings. The information in the table applies to both error and informational or warning messages.

Dialog Options for

Dialog Option	Description
Show message in popup dialog	Display message in a popup with an option to save the message to the clipboard. Selected by default. Clear the checkbox if you don't want system messages in a popup. Note: ESM also maintains system logs containing some audit information and details of any issues that occur. Refer to the Administration and Configuration section of the <i>ArcSight Command Center User Guide</i> and read the topic, "Log Retrieval."
Dialog Type:	
• Classic	Display the dialog in the front center of the ArcSight Console. The dialog remains in this position until you click OK to dismiss it.
• Animated	Animation defines the display duration, the dialog's direction of movement when it appears, and the direction of movement after the dialog times out.
	• Location: Position the dialog on one of the nine available locations on the screen and keep it displayed for the duration specified in Dialog Timeout. • For Entrance Animation: Dialog Timeout: Display the message in the number of milliseconds. The default is 3,000. Effect: For Fly, move the dialog from Direction and stops at Location. For Zoom, start the dialog at a small size and resize to its optimal size when it reaches Location. For Fade, make the dialog gradually appear at Location (ignore Direction). Direction: Move the dialog from one of eight origination points on the edge of the screen to Location. Direction works only with Fly and Zoom effects. Direction for Entrance Animation can be different from Exit Animation's.
	• For Exit Animation: Effect: For Fly, move the dialog from Location to Direction. For Zoom, shrink the dialog as it reaches the destination. For Fade, make the dialog gradually disappear at the same location when Dialog Timeout is reached (Direction is ignored). Direction: When Dialog Timeout is reached and if Effect is not Fade, move the dialog from Location to one of eight origination points on the edge of the screen. Direction works only with Fly and Zoom effects. Exit Animation and Entrance Animation can have different settings for Direction.

Setting Grid Options for the Viewer Panel

These options are for data displayed on the viewer panel's grid.

Where: Edit > Preferences > Grid View Options

Refer to the following table for available settings:

Grid View Options for

Grid View Option	Description
Font	Set global preference for font face, size, and style used in Grid views. Click into the Font field to get the drop-down menu arrow.  Click the arrow to bring up the Fonts dialog. Set the Font, Size and Style.
Color text by priority in grid	Apply distinguishing colors to the event rows in Viewer panel grid displays, based on their threat-priority levels. Note that this option can be overridden by the Color text by filter in grid option if conflicts occur. When these options are not selected, the text in grid rows defaults to black.
Color text by filter in grid	Apply distinguishing colors to the event rows in Viewer panel grid displays, based on the filters that selected them. You set these colors through the Configure button, described below. Note that this option, when selected, overrides the Color text by priority in grid option if conflicts occur. When these options are not selected, the text in grid rows defaults to black.
Pause the current channel on event selection	By default, selecting an event pauses the event flow to avoid scrolling. Clear this checkbox to allow the flow to continue regardless of a selection.
Do not prompt on verifying rule channel's timestamp change	Toggles on or off the option to have the system generate a prompt when the timestamp changes on an active channel populated by correlation events.
Do not prompt on channel restart	Toggles on or off the option to have the system generate a prompt when an active channel is restarted.
Check available database partitions on Active Channel start	This option applies to Oracle-based ESM and does not apply to ESM with CORR-Engine. If selected, this option causes the ArcSight Manager to recheck the status of available Oracle database partitions before starting an active channel. This does have a performance effect and is used only for certain historical analysis purposes.

Grid View Options for , continued

Grid View Option	Description
Default to 'Evaluate Once' time parameter for new Active Channels	Time evaluation parameter for new channel creation; default setting is Evaluate Once . You can choose Continuous as the time evaluation value.
Column Flip Limit	Determines the print format for Grid Views (channels, lists, and so forth). Grid views with the same or fewer columns than the Column Flip Limit print as a table, the same as is shown in the UI on the Console grid view. Grid views with more columns than the Column Flip Limit print details per row rather in a normal table like that shown on the Console grid view. The default setting for Column Flip Limit is "10" columns. (Tables with more than 10 columns print details per row.) See also "Printing from the Console" on page 93.
Filter Coloring Preferences	Click Configure to assign identifying colors to as many as five filters in the Configure Filter Colors dialog box.



Note: For instructions on customizing the grid's right-click option, **InActiveList**, see ["Customizing the Default Selections for Active Lists" below](#).

Customizing the Default Selections for Active Lists

If you are viewing events on an active channel, you have the ability to add selected events to existing active lists. By default, the Console's viewer panel enables you to browse to the resource locator so you can locate then select the desired list. These lists might be assigned to different list groups and might also be nested in a hierarchy.


If adding events from the event grid to existing lists is a frequent task for you, you can configure the grid's right-click option to display your top three frequently-used lists so that these lists are immediately available for selection.

Where: Edit > Preferences > Grid View Options


1. On the Grid View ActiveList Options area, click **Configure**. The ActiveLists resource selector is displayed.
2. Expand a group to locate your first preferred active list.
 - a. Select an active list and click **Add**.
 - b. Repeat to add up to a total of three lists.

- c. Change a list's position by clicking the up or down arrow.
- d. Remove lists from the selection as required.

Following is an example configuration for a selection of preferred active lists:



Following is the resulting default list selections when you open an event channel, right-click an event, and select **Active List > Add To**:



Note: This feature does not apply to the **Remove From** option from the grid view. If you are using the **Remove From** option, the Console displays an Active List selector dialog. You then navigate through the resource tree for active lists to select the list.

Setting Date and Time Formats

Purpose: Use the Date/Time option to choose a formatting style for the date and time strings displayed throughout the Console. You can also customize the details of any style you pick.

Where: Edit > Preferences > Date & Time

1. Click the **Formats** buttons and choose a date/time style from the lists for **Date & Time** Format and **Short Date & Time** Format options.
2. Select **Express all times as GMT** to universally show time values in GMT rather than local times.
3. Click **Apply** to put your changes into effect and leave the Preferences dialog box open, or **OK** to save your changes and close the dialog box.

If you want, you can customize the selected format string. Edit the **Format** string using the Java-style date options described in the **Format Help** window.


Setting Latitude and Longitude Options

Purpose: To define formats for latitude and longitude expressions in the **Asset > Locations** resource.

Where: **Edit > Preferences > Latitude & Longitude**

Choose from one of the available formats to express longitude and latitude.

Following is an example configuration for latitude and longitude format preferences:



Latitude & Longitude Entry Format

☒ DMS: Degree Minute Second (-4930'00")

☐ DM: Degree Minute (-4930.0')

☐ DD: Decimal Degree (-49.5000)


☒ Show compass direction (49.5000 S).

The options for latitude and longitude format vary from more exact to less so. Latitude and longitude can be shown in degrees, minutes, and seconds; degrees and minutes; or decimal degrees only. Additionally, an indicator of compass direction for the specified location can be shown or hidden in the editor.

To view the effects of your preference settings:

1. Choose **Assets** in the Navigator, click **Locations**
2. Create new location or edit an existing one to open up the **Location Editor**. (See ["Managing Locations" on page 172.](#))

Following is an example of how the Location Editor displays the preferred formats for Latitude and Longitude attributes:



Inspect/Edit

Location: Pacific

Attributes Notes

Location

* Name	Pacific
Latitude	0° 0' 0" S
Longitude	0° 0' 0" W
City	
Region Code	
Postal Code	
Country	N/A

Configuring Event Graphs

Purpose: You can modify the way graphs plot events, choosing to keep the source-event-target visual relationships compact; or to emphasize unique sources, targets; or both, in order to clarify the nature of attacks or situations.

Where: Edit > Preferences > Event Graph

Click the **Value** fields of the graph attributes to choose appropriate options:

- **Show Event Nodes:** Choose a basis for visually expanding or aggregating event nodes, relative to their source and target node instances.

Choice	Description
Once per common event	Graph only one instance of a given event node, regardless of the number of unique sources and targets that have it in common. For example, if sources 1 and 2 are directing the same event at targets 1, 2, and 3, there may be visual instances for each source and target, but only one of the event node.
Once per unique source	Graph one instance of a given event node per unique source, regardless of the commonality of associated targets. For example, if sources 1 and 2 are directing the same event at targets 1, 2, and 3, there are two visual instances of the event in support of the two distinct sources.
Once per unique target	Graph one instance of a given event node per unique target, regardless of the commonality of associated sources. For example, if sources 1 and 2 are directing the same event at targets 1, 2, and 3, there are three visual instances of the event in support of the three distinct targets.
Once per unique source or target	Graph one instance of a given event node per unique source-target pair, regardless of the commonality of the events involved. For example, if sources 1 and 2 are directing a given event at targets 1, 2, and 3; and as a chain, targets 1, 2, and 3 are sourcing the same events on to targets 4, 5, and 6; then there are six visual instances of the event in support of six distinct targets.

- **Show Source/Target IP Addresses as:** In cases where one source-event-target chains to another, you can choose to graph a source/target IP address as a single node, or to graph both the source and target instances of such an IP address.

Choice	Description
Distinct nodes	Visually plot both the source and target instances of a chained IP address.
Simple nodes	Visually plot a single node for an IP address that represents both source and target.

- **Source Node Identifier:** Choose a different event attribute to use as the identifier for source nodes. The default attribute is Source Address. Note that while all attributes are available, not all are appropriate choices for this purpose.
- **Event Node Identifier:** Choose a different event attribute to use as the identifier for event nodes. The default attribute is **ArcSight Category**. Note that while all attributes are available, not all are appropriate choices for this purpose.

- **Target Node Identifier:** Choose a different event attribute to use as the identifier for target nodes. The default attribute is **Target Address**. Note that while all attributes are available, not all are appropriate choices for this purpose.
- **Graph Layout:** Set the layout for all event graphs.



Note: You can override this default layout setting when you are actually viewing an event graph. For more details, refer to the topic, "Event Graphs as an Investigation and Analysis Tool" in *ESM 101*.

Hierarchical Layout	Display the event graph in tree-like nodes to show a related, sequential flow.
Organic Layout	The default layout.
Circular Layout	Display the source node as the center and the destination nodes arranged in a circle around the source.
Orthogonal Layout	Display the edges of the graph to run horizontally or vertically, parallel to the layout's X and Y axes.

- **Default Field Set:** Choose from the ArcSight-provided field sets to supply the data points in the graph. The default field set is from /All Field Sets/ArcSight System/Event Field Sets/Active Channels/Standard.

Setting Notification Popups

Purpose: You can manage received notifications from within the Console. In the Preferences dialog box, you can set a severity threshold for notification popups and optionally play a sound when notifications arrive.

Where: Edit > Preferences > Notifications

For the **Severity threshold for notification popup**, increase or decrease the integer value to a priority value that is based on the level at which you want to be alerted.

Select **Play a sound when a notification message is received** to also emit a sound when the alert threshold is met. Browse to the file of your preferred audio alert.

Managing Hot Keys

The Console provides schemas for configuring keyboard shortcuts to common actions. These schemas come with the Console:

- \$default
- Schemas for users



Tip: Keep these reminders in mind:

- Schemas for users other than administrators are listed only for users who have set up custom shortcuts on this Console under their own logins.
- Custom shortcuts are available only locally. See ["Sharing Custom Shortcut Schemas" on page 65](#) for more information.



Note: If this Console does not use UTF-8 encoding, refer to the "Installing ArcSight Console" section of the [ESM Installation Guide](#), and read the topic "Character Set Encoding" under "Installing the Console."

Schemas for users are all based on the \$default schema. That is, user schemas inherit all \$default schema shortcuts.

Where: **Edit > Preferences > Manage Hot Keys**

Under Available shortcut schemas, the schema in use shows as **"(active)"** next to its name.

You can define a keyboard shortcut for each listed command. Each command can have a different (or the same) keyboard shortcut depending on which schema you have selected.

Keyboard shortcuts are pre-defined for common commands. For example, the pre-defined keyboard shortcut for the Select All command (`edit.selectAll`) is Ctrl+A.

You cannot edit commands shown in red on the Preferences dialog: for example, `edit.delete`, `edit.redo`, `edit.cut`, `edit.copy`, `edit.paste`, and so forth. The flyover tooltips on these commands also indicates they are not editable.

There are many commands listed for which no shortcut is provided (for example, `file.new.Report`, `file.new.Rule`, `navigator.reports`, `navigator.queryViewers`, and so forth).

Adding Shortcuts for Frequently-Used Resources

This first task is not initiated on the **Edit > Preferences** dialog, but rather from various resource contexts in the Console. But the results of setting up shortcut keys on selected resources are shown on the **Edit > Preferences > Managing Hot Keys** dialog, as described here.

Where: **Navigator > <Resource>**

For example, choose **Active Channels** in the Navigator, and select an active channel such as `/All Active Channels/ArcSight Administration/System Events Last Hour`.


To add a shortcut to a resource:

1. Navigate to and select the resource for which you want to add a shortcut.
2. Right-click and choose **Manage Hot Keys** from the context menu to open the shortcut setup dialog for this resource.
3. Select the action you want to take with regard to the resource. Each resource has its own set of action, such as Edit <resource> and Show <resource>.
4. In the **Press new shortcut** field:
 - Optionally, click the button () to display a drop-down menu where you can set the type of shortcut to add (mouse, tab, and so forth) and set limits on keystrokes. For example, if you want to set the shortcut on this channel to Ctrl+C+H, change the keystroke limit from the default of 1, to **2** keystrokes.
 - Type the keyboard sequence you want to associate with the command.


If the keyboard sequence you typed is not in use, a light gray no conflicts message is shown in the Shortcuts currently used by field. For example, if you selected navigator.rules, placed the cursor in the Press new shortcut field, and typed **Ctrl+Alt+X**, you would get the no conflicts message.

If you type a sequence that is already used by another shortcut, you get a message in the Shortcuts currently used by field stating which resource is currently using the shortcut. For example, the default shortcut for navigator.rules is **Ctrl+Alt+L**. If you typed **Ctrl+Alt+R** in the Press new shortcut field, the message states that this sequence is already in use for navigator.reports.)

If you continue with the assignment, you get a prompt asking whether you want to remove the shortcut from the other resource and add it to this new one.
5. Click **Assign** to associate the shortcut with the resource.



6. Click **OK** to save your changes and close the dialog.
7. Confirm your setting by selecting **Edit > Preferences > Managing Hot Keys** dialog.
8. On the list of commands, locate the resource for which you created the shortcut. Resources are shown by their URIs.
9. Select the URI to display the associated shortcut, as in the following example:



Modifying a Custom Shortcut


Shortcuts are associated with schemas based on the user.

Where: Edit > Preferences > Manage Hot Keys

To modify a custom shortcut:


1. On the Edit Preferences > Manage Hotkeys dialog, select a shortcut schema (the associated user) in which you want to modify shortcuts for commands.

In this example, the schema for the user called **admin** is selected. Note, however, that the schema selected for modifying a hot key need not be the “active” schema; as it happens to be in this example.



2. Select the command for which you want to modify the hot key.

You can filter for commands containing a given string (for example, **navigator** to find all navigator commands).



3. In the **Press new shortcut** field:

- Optionally click the button () to display a drop-down menu where you can set the type of shortcut to add (mouse, tab, and so on) and limits on keystrokes. The default keystroke limit is **1**. If you set it to **2** or **3**, you have more combinations of keystrokes available to use for custom settings.
- Enter the keyboard sequence you want to associate with the command.

If the keyboard sequence you entered is not in use, a light gray no conflicts message is shown in the Shortcuts currently used by field. For example, if you select `navigator.rules`, place the cursor in the Press new shortcut field, and press **Ctrl+Alt+X**, you get the no conflicts message.

If you enter a sequence that is already used by another shortcut, you get a message in the Shortcuts currently used by field telling you which resource is currently using the shortcut. For example, the default shortcut for `navigator.rules` is **Ctrl+Alt+L**. If you enter **Ctrl+Alt+R** in the Press new shortcut field, you get a message noting that this sequence is already in use for `navigator.reports`.)

If you continue with the assignment, you see a prompt asking whether you want to remove the shortcut from the other resource and add it to this new one.

4. Click **Assign** to apply the new shortcut to the command.



Tip: An asterisk is displayed next to commands for which the pre-defined shortcuts have been modified or overwritten. These customized commands are also displayed in blue text, rather than the usual black.



5. Click **Apply** or **OK**.

To modify a custom shortcut directly from the resource:

You can modify a custom shortcut for a resource in either of these ways:

- Directly from the right-click **Manage Hot Keys** dialog on that resource
 - From the **Edit > Preferences > Manage Hot Keys** dialog as described above
1. Navigate to and select the resource from which you want to remove the shortcut.
 2. With the appropriate resource selected, right-click and choose **Manage Hot Keys** from the context menu to bring up the shortcut setup dialog for this resource.
 3. Select the action (for example, Show or Edit) associated with the shortcut.
The shortcut is shown in the Press new shortcut field.
 4. Modify it as needed. See the previous procedure.
 5. Click **OK** to save your changes and close the dialog.

Removing a Custom Shortcut

Where: Edit > Preferences > Manage Hot Keys

To remove a custom shortcut (key sequence) for any command:

1. Select the schema in which you want to modify the command.
2. Select the command for which you want to modify the hot key.
3. Select one of the customized commands (customized commands are shown in blue text with an asterisk).



The current key sequence associated with this command is shown in the **Shortcuts for selected command** field.

4. Click the **Remove** button next to the Shortcuts for selected command field.

The custom shortcut (key sequence) is removed, and replaced by the default key sequence (if there was one).



Caution: As soon as you remove the shortcut by clicking **Remove**, the changes are saved. Even if you click Cancel to close the Preferences dialog at this point, *the original shortcut is not saved.*

For example, if `navigator.rules` was modified to be associated with `Ctrl+Alt+X`, then after you remove this shortcut, `navigator.rules` would again be associated with its default shortcut of `Ctrl+Alt+L`.



Tip: You can only remove custom shortcuts, but not default shortcuts.

To remove a custom shortcut directly from the resource:

You can remove a custom shortcut for a resource in either of these ways:

- Directly from the right-click **Manage Hot Keys** dialog on that resource
 - From the **Edit > Preferences > Manage Hot Keys** dialog as described above.
1. Navigate to and select the resource from which you want to remove the shortcut.
 2. With the appropriate resource selected, right-click and choose **Manage Hot Keys** from the context menu to bring up the shortcut setup dialog for this resource.
 3. Select the action (for example, **Show** or **Edit**) associated with the shortcut.
The shortcut, if any, is shown in the Press new shortcut field.
 4. Click **Remove**.
 5. Click **OK** or **Cancel** to close the dialog.



Caution: As soon as you remove the shortcut by clicking **Remove**, the changes are saved. Even if you click **Cancel** to close the Preferences dialog at this point, *the original shortcut is not saved*.

Activating a New Shortcut Schema

For more information on schemas, see the introduction to the shortcut key management at ["Managing Hot Keys" on page 58](#).

Where: **Edit > Preferences > Manage Hot Keys**

To activate a new schema:

1. Select the schema you want to activate.
2. Click **Set Active**.



Tip: To get an enabled Set Active button, select a schema that is not currently applied. If you select a schema that is already active, the Set Active button is disabled.

3. Click **Apply** to apply the new schema, or click **OK** to apply the new schema and close the Preferences dialog.

Sharing Custom Shortcut Schemas

Shortcut schemas are available only to the local Console. That is, if schemas for several different users are configured on a Console running on a particular machine, those shortcut setups (schemas) are not available for the same Console user logins on other machines.

This means that if you want the same shortcuts to exist in other Console installations, you must manually set these up in those installations.

Viewing

This section provides information on using the Console Viewer Panel and choosing look-and-feel options (skins) for the .

The Viewer Panel

You see the results of security-event analyses in the Viewer panel, which can display several different types of views. (See also ["Using Views" on page 231.](#))

Although there are some views that display information about resources, most views are active channels, which are continuously evaluated collections of security-event data. (See also ["Monitoring Active Channels" on page 224.](#))




Tip: Here are some Viewer Panel features you can use.

- To show a resource (like a particular dashboard or active channel) in the viewer, right-click it in the Navigator tree and select **Show <resource>**.
- To close individual views quickly, **Shift+click** their name tabs. (You can also right-click a view's name tab and select **Close** from the popup menu.)
- To float the Viewer panel, click the **Float** icon at the top left of the Viewer.

The **Viewer** tabs in the Viewer panel have a live link at the top. You can click these links to open the contents in an external, fully functional browser window.

For security reasons, HTML that might include JavaScript, plug-ins, or other embedded objects are rendered in the default browser you specify through the Preferences dialog box. The default browser is also used by PDF document files.



If your Console is not already displaying a default set of pre-defined views, or if you want to change the views displayed, you can use these options:

- Choose **Window > Viewer Panel** to open the panel if it isn't open.
- Choose the **Active Channels**, **Dashboards**, or **Threat Detector** resource trees in the Navigator panel to find analysis tools or results to view.
- Right-click a resource in a tree and choose **Show <resource>** to open it in the Viewer panel.
- When multiple tabbed views are open in the panel, click the tabs at the **top** of the panel to choose the active channel you want to see, and the tabs at the **bottom** of the panel to choose which view of that active channel should be foremost.

To close an individual view, **Shift+click** its name tab. (You can also right-click a view name tab and choose **Close** from the popup menu.)

Using active channels and the many types of views they offer is fully covered in the topics under these headings:

- ["Monitoring Events" on page 224](#)
- ["Selecting and Investigating Events in Active Channels" on page 291](#)
- ["Using Dashboards" on page 252](#)

Look-and-Feel

If you start the from the command line with the **arcsight console** command (in ARCSIGHT_HOME/current/bin), use the **-laf <style>** flag to specify a look-and-feel style. For example, the following command starts the Console with a “metal” look-and-feel:

```
arcsight console -laf metal
```

The other possible styles are plastic, the default for Unix, and plastic3d.

These styles modify the Console display and associated online help.

The screen captures and illustrations used throughout the online help show various look-and-feel styles.

Inspecting and Editing

ArcSight Console provides the Inspect/Edit panel to examine the details of events that appear in active channels in the Viewer panel, or to modify the resource attributes in the Navigator panel. You can examine security events through the Inspect/Edit panel's Event Inspector, and edit resources using specialized editors, one for each specific resource type.




Note: Press **Enter** to register edits made in editors and channel columns.

To ensure that ESM registers a change you make to a field in editor and channel columns, press Enter before clicking **Apply** or **OK**.

Overview of Inspect/Edit Features and Utilities

Each editor has its own controls and attributes, described in the Help for each resource.




The Inspect/Edit panel opens automatically if you double-click an event in a grid view or choose to edit a resource in the Navigator panel. You can also right-click an event in a grid view and choose **Show Event Details**.

In the Inspect/Edit panel, you can:

- Choose **Window > Inspect/Edit Panel** to open or restore the panel, if it already has inspectors or editors in it. If no inspectors or editors are open, the panel does not display anything.
- If no editors or inspectors are open, or to work with different ones, double-click an event in a grid view or right-click an item in a Navigator panel resource tree and choose **Show <resource>**.

- To clear an editor from the Inspect/Edit panel, right-click its tab and choose **Close**.
- Click the **Field Set Selector** dropdown menu (defaults to Event Inspector at Console startup) to use your field set of interest.




Note: If you have not exited the Console for a day or more, you may notice that the Field Set Selector no longer displays a list of available field sets. If this happens, right-click on any field under the Name column and choose **Select a Field Set**. The dropdown works with newly-started Consoles. As a good practice, exit the Console if you are done using it.

- Click the **Hide Empty Rows** button () to see only populated fields.
- Click the **New Field Set** button () to create a new field set.
- Click the icon toggle button () to show/hide icons next to each field entry.

Searching for Fields in Event Inspector, Resource Editors, or CCE

To find an item in a list of fields on the Event Inspector, any Resource Editor, or the [Common Conditions Editor \(CCE\)](#), start entering the search string in the Search for field at the bottom of the panel. The search is predictive in that it will navigate to and select matching fields as you type. The Search utility works essentially the same way in the Event Inspector and in resource editors that use field sets and filters (and, by association, the CCE).



If you start to type a term that is not in the field list, the search text turns red. If you backspace and start deleting text, the text will change from red to black when a matching field is found. Resume typing to find another matching term.

To exit the Search, press the **Enter** key.

Getting More Help

The best way to learn more about the Event Inspector and each of the many resource editors is to click the question mark button (?) in the upper-right corner of the Inspect/Edit panel or the **Help** button () in the lower right of a resource editor.

Controlling the

The has certain common controls for basic tasks like copying and pasting, and showing or hiding panels or the status bar.

There are four toolbars under the Console menus. Each button has an identifying tool tip, but the full descriptions are as follows.


To show or hide toolbar components, right-click the toolbar and select or deselect the sections you want to change.

's Toolbar Components

Command Group	Icons	Functions
File		New resource, Open, and Save. Saving and opening applies to settings (.ast) files.
Edit		The Cut, Copy, Paste, Delete, and Search , buttons operate as they do in any application. Cutting, copying, and pasting applies to text and resources.
Channel controls		The Replay buttons have the same functions in certain views on the Viewer panel as their counterparts do on VCRs or CD players. From left to right, the buttons are: Rewind to Start, Rewind Incrementally, Pause, Play, Stop, Go Forward Incrementally, and Go Forward to End. Use the Replay buttons when working with channels configured for this mode.
View		The Notifications button, if blue, indicates there are no new notifications. The button turns red if you have messages to acknowledge. Click the button to open the Notifications manager in the Viewer panel so you can acknowledge the notification and resolve the issue. The Slide Show button enables you to start an image dashboard slide show and set the interval.
Window		Click the Show/Hide buttons to open or close the Navigator, Viewer, and Inspect/Edit panels; and status or menu bars. Click the Floating button to bring floating windows forward.
Network Tools		These buttons run standard IP-based network analysis tools as described in "Using the Network Tools" on the next page.
System		Open a scheduled jobs list and add user categorizations to selected events .
Status Bar		The status bar is across the bottom of the Console window. Use the Window > Status Bar option to toggle the bar on or off. When the status bar is showing, it displays Console operation messages. Normal status messages appear in blue and error messages are in red. To view details on a message, click the message in the status bar. The ArcSight Messages dialog is displayed with the current message highlighted. From this dialog, you can access Console messages, system messages and user notifications. To copy any message from the Messages dialog, highlight it and click Copy. The message is copied to the clipboard along with associated date and time. You can then paste the message into any other window, mail program, or editor that accepts ASCII text.

To save error and warning messages:

While using the Console interface, certain error messages, warnings, and notifications may appear in a small dialog box:



To capture the message and supporting data, click the **Copy** button or check **Copy message to system clipboard** to copy the entire message to the Clipboard. You can then paste the error message in text fields in the ArcSight Console, into the body of an e-mail message, or other applications.

Using the Network Tools

The network tools shown here are also available from the Tools menu:



ArcSight provides **Network Model**, **Use Case**, **Name Server Lookup**, **Ping**, **Port Usage**, **Trace**, **WebSearch**, and **Who is** as default utilities. Most of these tools are utilities you use to investigate events in grid views. In a grid view, you right-click an event to access these tools from a context menu. A wizard-based utility called **Send Logs** gathers logs and diagnostic information for review or which you can email to customer support.

You can add, copy, edit, or delete network tools using the Tools menu in the menu bar. The toolbar buttons and menu commands adjust automatically to such changes.



Tip: The Network Tools are also available as *integration commands* (see ["Network Tools as Integration Commands" on page 692](#)).

These tools are available in both places on the Console UI, but for future releases the legacy "network tools" feature described here will be phased out in favor of the integrations commands. The same, customizable tools and commands will be available (**ping**, **whois**, and so on), along with other new commands and a full set of application integration features.

To configure these tools, choose menu option **Tools > Local Commands > Configure**, as described in the following topics:

- ["Running a Tools Command" on the next page](#)
- ["Adding or Editing a Tool" on page 75](#)

Running a Tools Command

To run a tools command:

1. On a grid view, select an IP address.
2. Right-click and select **Tools**, then select one of the tool options described here:

Tool Options

Tree	Icon	Resource
Network Model		Configure the network model. This button launches the Network Model wizard.
Use Case		Configure a use case. Instructions are in the documentation that comes with each optional Security Use Case.
Send Logs		Access this from the Tools > Send Logs menu. Start the Send Logs wizard to gather logs and diagnostic information. Logs and diagnostics can be collected for all or a selected set of ArcSight components. (See "Send Logs" on page 1047.)
Local Commands:		
Nslookup (Windows)		Resolve an IP address to a host or domain name, or vice versa. Run the command on the Console on Windows.
Nslookup (Linux)		Resolve an IPv4 address to a host or domain name, or vice versa. Run the command on the Console on Linux.
Nslookup-IPv6 (Linux)		Resolve an IPv6 address to a host or domain name or vice versa. Run the command on the Console on Linux.
Ping (Windows)		Determine if an IP address in the selected cell is reachable on the network. Test and debug a network by sending a packet and waiting for a response.
Ping (Linux)		Determine if an IPv4 address in the selected cell is reachable on the network. Test and debug a network by sending a packet and waiting for a response. Run the command from a Console on Linux.
Ping6 (Linux)		Determine if an IPv6 address in the selected cell is reachable on the network. Test and debug a network by sending a packet and waiting for a response. Run the command from a Console on Linux.
Portinfo (Windows)		List standard usage, for example, WWW, FTP, and so on for a specified port number.
Portinfo (Linux)		Find information about the selected port. Run the command from a Console on Linux.
Traceroute (Windows)		Show the path from the to the IP address selected in the grid view, reporting the IP addresses of all routers in between.

Tool Options, continued

Tree	Icon	Resource
Traceroute (Linux)		Show the path taken by packets across an IP network. Run from a Console on Linux.
WebSearch		Search the Web through Google to find links to the keywords present in currently selected active channel grid view cells.
Whois (Windows)		Look up the owner of a given domain name or an IP address; information might include addresses and telephone numbers.
Whois (Linux)		Look up the owner of a given domain name or an IP address; information might include addresses and telephone numbers. Run the command from a Console on Linux.

3. Based on the tool selected, a window appears with the information.
4. In the window, click **Close**.

Adding or Editing a Tool

To add or configure (edit) a tool:

1. Choose the menu command **Tools > Local Commands > Configure**.
2. In the Configure Tools window:
 - Click **New** if you are adding a tool, or
 - Select an existing tool and click **Edit**.
3. In the Tool window, set options for command line parameters to be used for the program, described below:

Tool Configuration Options

Option	Description
Name	User-friendly name for this tool.
Program	Path to the executable file.
Working Directory	Default location assumed for arguments to the command. For example, to create a command (for example, delete <file>.ast) that acts on a file type that always resides in the same directory, specify the location here to save users from having to provide the full path to the file each time they use the command.
Icon	Path to the icon image file used to represent the tool.

Tool Configuration Options, continued

Option	Description
Program Parameters	Provide any parameters needed for the command. You can enter parameters in the field, or click the  button to get a pull-out menu where you can select Event Attributes to use as parameters, or add the selected cell or selected row as parameters to the command.
Show in toolbar	When Show in toolbar is on, the tool icon is shown in the Console toolbar. By default, this option is selected.
Use with data export	The purpose of this option is to separate tools that are run against events in channels and tools used as a destinations for event export. By default, this option is not selected (off). If this tool is to be used as a destination for event export, select Use with data export. If this tool contains a command that will run against events in a channel, leave Use with data export off.

4. **Name, Program, Working Directory, Icon, and Program Parameters** (command line parameters to be used for the program) are text fields. Also select whether you want the tool to show in the toolbar
5. Click **OK**, then **Done**.

To delete a tool:

1. Choose menu command **Tools >Local Commands > Configure**.
2. In the Configure Tools window, select an existing tool and click **Delete**.
3. In the dialog box, click **Yes**.
4. Click **Done**.

Staying Informed

This topic discusses ways by which the Console helps you stay informed about developing situations involving events, and critical system status.

Topics include:

- ["Acknowledging Notifications" on the next page](#)
- ["Checking the Status of the Distributed Correlation Cluster" on the next page](#)
- ["Using Notes" on page 79](#)
- ["License Tracking" on page 80](#)

Acknowledging Notifications

To be informed when certain defined events or circumstances occur. You might receive notifications by cell phone message or e-mail but you can be sure to see an indicator in the **Notifications** button on the Console toolbar:



Notifications can be sent as a result of a rule action, or by another user monitoring events in a grid. Clearing a notification requires that you acknowledge it. Whether or not you need to take other action depends on the circumstances. Acknowledgments are described briefly here, but for full detail, see ["Managing Notifications" on page 216](#).

To acknowledge a cell phone message:

Acknowledge a call by replying to the e-mail sent through your cell phone. An e-mail enabled cell phone is required for receiving notifications and replying to them.

To acknowledge an e-mail message:

Acknowledge an e-mail message by replying to the message. Reply to the e-mail address from which the notification was sent.

To acknowledge notifications at the :

The automatically alerts you of pending acknowledgments. The **Acknowledge Notifications** button is automatically enabled when you have one or more notification messages to be acknowledged. When you click the **Acknowledge Notifications** button, the Notifications manager opens in the Viewer panel so you can acknowledge and resolve the notification.

Checking the Status of the Distributed Correlation Cluster

This topic applies to ESM installed in **distributed correlation** mode.

The Console toolbar contains the **Cluster View** icon to show the health of your distributed correlation cluster. It provides the link to the Cluster View dashboard on the ArcSight Command Center.

Cluster Status Color Indicators

Color	Meaning
Red 	Indicates any of these conditions: • All aggregators are down. • All correlators are down. • All connections to MBus are down. • All connections to DCache are down. • If the properties are set, the message lag on correlator or aggregator is above threshold set for a Red icon. See the topic, Defining Message Lag Thresholds for procedures.
Yellow 	Indicates any of these conditions: • An aggregator is down. • A correlator is down. • Some connections to MBus are down. • Some connections to DCache are down. • If the properties are set, the message lag on correlator or aggregator is above threshold set for a Yellow icon. See the topic, Defining Message Lag Thresholds for procedures.
Green 	Cluster is operational.

To access the Cluster View dashboard on the Command Center:

1. Click the **Cluster View** icon.
This launches the login popup to the ArcSight Command Center.
2. Enter your login credentials and click **OK**.

Defining Message Lag Thresholds

You can define two levels of thresholds each for the aggregator and correlator: thresholds to change the Cluster View icon to yellow, and thresholds to change the icon to red. By default, no message lag thresholds are tracked through the icon colors.

The following table describes property settings to include in the properties file:

Property	Description
<code>aggregator.lag.alert.yellow.threshold</code>	The threshold of message lag at the aggregator. Specify a positive integer value. If message lag is above your specified value, the Cluster View icon turns to yellow.
<code>aggregator.lag.alert.red.threshold</code>	The threshold of message lag at the aggregator. Specify a positive integer value. If message lag is above your specified value, the Cluster View icon turns to red at the aggregator.
<code>correlator.lag.alert.yellow.threshold</code>	The threshold of message lag at the correlator. Specify a positive integer value. If message lag is above your specified value, the Cluster View icon turns to yellow.
<code>correlator.lag.alert.red.threshold</code>	The threshold of message lag at the correlator. Specify a positive integer value. If message lag is above your specified value, the Cluster View icon turns to red.

To set the threshold values for the message lag:

See "Managing and Changing Properties File Settings" in the [ESM Administrator's Guide](#) for the proper way to edit the `server.properties` file.

To the `server.properties` file, add the properties mentioned in the above table to the thresholds you want. If you enter `-1`, then the Cluster View icon will not be affected by any message lags in aggregators and correlators.



Note: The color of the Cluster View icon is affected by a combination of additional factors, like MBus and DCache connections, in addition to message lag in aggregators and correlators if you set the thresholds.

Using Notes

Each resource and resource group in the Navigator panel has an editor, and each editor has a Notes tab. The Notes tab retains all the text that you and others add to the resource.

Notes tabs have Table and List sub-tabs to show you tabular or text layouts of the notes accumulated for a resource. Notes are stored chronologically and you can sort them by clicking the **Date**, **Owner**, and **Text** headers.

To add a note:

1. On the Navigator panel resource tree, right-click a resource group or individual resource.
2. For a resource group, choose **Edit Group**. For a resource, choose **Edit <resource>**.
3. In the Inspect/Edit panel, click the editor's **Notes** tab.
4. In the Notes space, type a note.

5. Click **Save** and then **OK**.

To view a note:

1. On the Navigator panel resource tree, right-click a resource group or individual resource.
2. For a resource group, choose **Edit Group**. For a resource, choose **Edit <resource>**.
3. In the Inspect/Edit panel, click the editor's **Notes** tab.

To delete a note:

1. On the Navigator panel resource tree, right-click a resource group or individual resource.
2. For a resource group, choose **Edit Group**. For a resource, choose **Edit <resource>**.
3. In the Inspect/Edit panel, click the editor's **Notes** tab.
4. Right-click a note and choose **Delete**.

To search for text strings in Notes:

You can run a search on a resource's Notes tab. Refer to the topic, ["Finding Resources" on page 704](#). That topic provides instructions on using the Search field on the Console's toolbar and entering correct search syntax.

License Tracking

The product tracks the status of licenses for features you use, including actors, user limits, Command Center user limits, device number limit, actor and asset number limit, and events-per-second limit.

Licenses for the features available to you are installed and configured at setup time. For details about setting up licenses during the installation and configuration process, see the [ESM Installation Guide](#).




Note: License tracking includes disabled and deleted actors.

The ESM license tracking feature includes actors that are still in the ESM actor model with the status `Disabled` or `Deleted` in IDM. ESM's identity management feature preserves disabled and deleted actors in the actor model to track any unauthorized activity related to disabled or deleted actors.

If you do not want the ESM license tracking feature to evaluate actors with the status `Disabled` or `Deleted` in IDM, you can manually remove them from the ESM actor model. However, manually removing disabled or deleted actors also removes the ability for ESM to track unauthorized activity related to these accounts. For details, see ["Deleting Actors" on page 793](#).

License Tracking Notifications

If your feature usage is close to or has exceeded the license agreements for your organization, you see a notification dialog when starting up the , for example:



Your access to these features remains in place even if the license limit has been exceeded.

Standard Reports for License Status Tracking

You can check on the status of your ESM feature licenses using the reports and focused reports found under All Reports/ArcSight Administration/ESM/Licensing.

Each report is described in the topic "ESM Licensing" in the [ArcSight Administration and ArcSight System Standard Content Guide](#).

For details about running reports, see ["Running and Managing Reports" on page 446](#).

Using the File Menu

Keyboard shortcut: **Alt+F**

See also ["Keyboard Shortcuts \(Hot Keys\)" on page 89](#).

Options on the File Menu

Option	Icon	Resource	Shortcut
New		Create a new resource from the available submenu.	
Open		Open an existing Console settings file to use that configuration.	Ctrl-O
Save		Save your latest Console settings in the current configuration (.ast) file.	Ctrl-S
Save As		Save your current Console settings in a different configuration (.ast) file.	
Save to Manager		Save your current Console settings at the ArcSight Manager rather than locally, so you can get these settings at a different Console.	

Options on the File Menu, continued

Option	Icon	Resource	Shortcut
Load From Manager		Load a preferred Console configuration file (.ast) from the ArcSight Manager, so you can use it with this Console.	
Send To		Send a local Console configuration (.ast) file to an e-mail address so another user can save and use it with their Console.	
Log Out		Log out of the Console with your current user ID, without exiting, so someone else can log in.	
Exit		Log out of the Console and exit.	Alt-F4

Using the Edit Menu

Keyboard shortcut: **Alt+E**

See also ["Keyboard Shortcuts \(Hot Keys\)" on page 89](#).

Options on the Edit Menu

Option	Icon	Resource	Shortcut
Cut		Cut selected text.	Ctrl-X
Copy		Copy selected text or resources.	Ctrl-C
Paste		Paste text or resources from the clipboard.	Ctrl-V
Delete		Delete selected text or resources.	Delete
Select All		Select all text.	Ctrl-A
Preferences		Open the Preferences dialog box to make personal configuration changes.	
Find Resource		Use the Find Resource query editor to search for resources and review their details.	Ctrl-F

Using the View Menu

Keyboard shortcut: **Alt+V**

See also ["Keyboard Shortcuts \(Hot Keys\)" on page 89](#).

Options on the View Menu

Option	Icon	Resource	
New Active Channel		Open the New Active Channel dialog box so you can set up and start a new active channel in the Viewer panel.	Ctrl+Shift-D
Show Active Channel		Open the Active Channel Selector dialog box so you can choose an active channel to display in the Viewer panel.	Ctrl+Shift-S
Recent Active Channels		Choose a recently opened active channel to display in the Viewer panel again, if available.	
Resource Hotkeys		Show currently programmed keyboard shortcuts for actions on the Console. These keyboard shortcuts are defined in the Console Preferences dialog (Edit > Preferences > Manage Hotkeys). For more information, see "Keyboard Shortcuts (Hot Keys)" on page 89 .	
Show my group's cases		Display a Case Channel showing cases assigned to your user group.	
Show a group's cases		Display a User Selector for user groups, from which you expand and select the group you want. Click OK to display the Case Channel. If no cases are assigned to the group you selected, the channel is empty and will state that no data matches your query.	
New Dashboard		Create a new, untitled and empty dashboard to populate with data monitors.	Ctrl+Shift-B
Show Dashboard		Open the Load Dashboards dialog box so you can select dashboards to open in the Viewer panel.	Ctrl+Shift-W
Recent Dashboards		Choose a recently opened dashboard to display in the Viewer panel again, if available.	
Themes		Choose between Default and Dark. See "Changing the Console Display" on page 47 for descriptions.	
Notification Acknowledgement		Show all Notifications for the current user (pending, undeliverable, not acknowledged, acknowledged, and resolved)	Ctrl-N
Show Messages		Show all Console messages, system messages, and user notifications in the ArcSight Messages dialog.	Ctrl-M
Next View		Take you to the next open view or tab in the Viewer panel.	Ctrl+Shift-N
Previous View		Take you to the previous open view in the Viewer panel.	Ctrl+Shift-P
Close All Views		Close all views that are open in the Viewer panel.	
Slide Show		Show a continuous slide show of all open channels and dashboards.	F11 (Toggle to start or stop)

Using the Window Menu

Keyboard shortcut: **Alt+W**

See also ["Keyboard Shortcuts \(Hot Keys\)" on page 89.](#)

Options on the Window Menu

Option	Icon	Resource	Shortcut
Navigator Panel		Show or hide the Navigator panel.	Ctrl-1
Viewer Panel		Show or hide the Viewer panel.	Ctrl-2
Inspect/Edit Panel		Show or hide the Inspect/Edit panel.	Ctrl-3
Status Bar		Show or hide the status bar.	Ctrl-4
Floating		Bring to the front one of the listed floating (undocked) windows, if available. Disabled if all viewer panels are docked.	

Using the Tools Menu

Keyboard shortcut: **Alt+T**

See also ["Keyboard Shortcuts \(Hot Keys\)" on page 89.](#)

Options on the Tools Menu

Option	Sub-menu	Icon	Resource	Shortcut
Local Commands	Configure		Add, copy, edit, or delete Network Tools.	Alt-C
	Results		Display the Tool Results dialog box.	Ctrl+Shift-R
	Nslookup		Resolve an IP address to a host name.	
	Ping		Determine whether an IP address is online.	
	PortInfo		List the default protocol usage for a specified port number (for example, WWW, FTP, SMTP).	
	Traceroute		Show the path to an IP address.	
	WebSearch		Use Google to search the web for event-related keywords.	
	Whois		Find the registered owner of a given domain name.	

Options on the Tools Menu, continued

Option	Sub-menu	Icon	Resource	Shortcut
Network Model			Launch the Network Model wizard. See "Populating the Network Model Using the Wizard" on page 132.	
Use Case			Launch the Use Case wizard. Refer to the specific use case document or instructions. See also "Use Cases" on page 41.	
Send Logs			Launch the Send Logs wizard. See "Send Logs" on page 1047.	

Using the System Menu

Keyboard shortcut: **Alt+S**

See also ["Keyboard Shortcuts \(Hot Keys\)" on page 89.](#)

Options on the System Menu

Option	Icon	Resource
Scheduled Jobs		Open the Job Scheduler. For more information, see "Job Scheduler" on page 1003.
Categorize Event		Select a non ArcSight event in the grid, then select System > Categorize Event menu option to apply a category.

Using the Help Menu

Keyboard shortcut: **Alt+H**

See also ["Keyboard Shortcuts \(Hot Keys\)" on page 89.](#)

Options on the Help Menu

Option	Icon	Resource	Shortcut
Browse Documentation		Open the index page of the embedded ESM documentation set in HTML.	F1
What's New		Open the Online Help's "What's New" topic.	
Console Guide		Open the Online Help's landing page with links to major topics.	
Software Support Online		Open a browser window that displays the Software Support home page, so you can sign in and access the case manager, downloads, communities, and other features.	
About		Show your ArcSight installation's legal notices and version information.	

Using Right-Click Context Menus

Right-click context menus appear throughout the Console. This section describes common options available from right-click menus in different contexts. Context menus in different resources can offer other options specific to that resource. To understand all the options available for a particular resource, see the topic related to that resource.

The Navigator panel presents individual resources and groups that help organize them. Here are the common options available from the right-click context menus in the Navigator panel. Not all options are available in all contexts; those that are not available will appear in grey text. The details of many of these options, such as creating a new resource, are described in the topics dedicated to that resource.

Common Right-Click Menu Options on the Navigator Panel

Option	Applies to	Description
New <Resource>	Resources	Open the editor for the selected resource to allow you to create a new one.
Edit <Resource>	Resources	Open the editor for the selected resource to allow you to edit an existing one.
Delete <Resource>	Resources	Initiate a delete sequence for the selected resource. A confirmation step is required before the resource is permanently deleted.
Add to Favorites	Active Channels, Actors, Assets, Cases, Connectors, Customers, Locations, Networks, Vulnerabilities	Add the selected resource to the Favorites list.
Find in Resource Tree	Resources in Recents and Favorites	From the Recents or Favorites list, locate and select the resource in the resource navigator panel.
Remove from <Recent or Favorites>	Resources in Recents and Favorites	Remove the resource from the Recents or Favorites list.
- Recon - Recon (Multiple Fields)	Event Channels Enabled if ESM is configured to connect to a Recon deployment.	- Execute a search on Recon for the selected value. - Execute a search on Recon for values of up to five fields. See "Running Recon Searches" on page 293.

Common Right-Click Menu Options on the Navigator Panel, continued


Option	Applies to	Description
Integration Commands	Resources and Channels	From Console, link to other ArcSight applications and tools. For more information, see "Integration Commands " on page 664.
Add to Package	Resources	Add the selected resource to an existing package. For more about packages, see "Managing Packages" on page 730.
Show <Resource>	Resources	Display results gathered by the resource in the Viewer panel.
Debug Event Priority	Event Channels	Display an event's priority information, which includes scores for severity, relevance, model confidence, and asset criticality. For more information, see "Priority Calculations and Ratings" on page 1019.
Graph View	Resources	Display a graphical view of the resource in relation to other associated resources. For more about resource graphs, see "Visualizing Resources" on page 714.
New Group	Groups and Resources	Add a new group. The group's attributes are defined in the Inspect/Edit panel.
Edit Group	Groups and Resources	Edit an existing group through the Inspect/Edit panel. You can edit a variety of group attributes such as the name, description, owner, and so on. Available in all resources.
Delete Group	Groups and Resources	Delete a group.
Rename	Groups and Resources	Change a group's or single resource's name directly on the Navigator pane without going through the group's Inspect/Edit panel. Caution: Be careful about renaming a resource which has, or which will eventually have, dependent resources. Once you change the name, don't re-use the old name for a new resource of the same type because the dependent resources may continue to refer to the new resource with the old name. An example of a dependent relationship is of a trend with queries that depend on that trend.
Edit Access Control	Groups and Resources	Launch the Access Control Editor. For more about the Access Control Editor, see "Editing Access Control Lists (ACLs)" on page 107.
Show Invalid Reason	Groups and Resources	For a group or resource shown as invalid (improperly constructed), display the explanation for the invalidity.
Validate <group or resource>	Groups and Resources	Validate the group or resource that is shown to be invalid because the group or resource was not constructed properly. For more information, see "Validating Resources" on page 718.
Lock <group or resource>	Groups and Resources	Prevent a group or resource from being edited by users other than the creator of the information.
Unlock <group or resource>	Groups and Resources	Allow edits to the group or resource.

Common Right-Click Menu Options on the Navigator Panel, continued

Option	Applies to	Description
Set deprecated flag	Groups and Resources	Set the Deprecated check box on the group or resource's Attributes tab as seen in the Inspect/Edit panel.
Remove deprecated flag	Groups and Resources	Remove flag (clear the Deprecated check box) from a previously-deprecated group or resource.
Refresh	All	Updates the Console with the latest changes.
Knowledge Base	Groups and Resources	• Show Article—Display the associated Knowledge Base article for the resource or group. The associated article is displayed in the Viewer panel. • Associate With—Select a Knowledge Base article from the Knowledge Base Article Selector to be associated with the resource or group. • Association Help—Display the Help window for associating a resource with a knowledge base article. For more information about knowledge bases, see "Knowledge Base Authoring" on page 698.
Reference Pages	Groups and Resources; certain events	Display pointers to additional reference information, if such information is available for the group or resource. For more information, see "Reference Pages" on page 1027.
Print <resource> Tree		Print a selected resource's tree view. For more about using this printing feature, see "Printing Navigation Tree Views of Resources" on page 93.
Create Channel with filter	Customers, Connectors, Rules, Vulnerabilities	Create a channel of the selected resource. When used with a rule (standard rule only), the channel shows audit events generated for that rule. If the channel is empty, this means the rule has not triggered.
Show Assets	Locations, Vulnerabilities	Displays assets with the selected location or vulnerability
Vulnerability Assets Report	Vulnerabilities	Creates a report of assets with vulnerabilities
Help	Groups and Resources	Launch Console Help topic for the selected resource.

Using the Advanced Selector While Editing Resources

Some resources need a resource attribute. For example, a query viewer needs a query to get data from the database. The **Advanced Selector** button on the Edit panel of some resources provides options to search for, then select a resource.




Clicking the button displays the Advanced Selector dialog.


Search is enabled by default.

1. If you know the resource by name, enter as many characters in the text field to match the resource name. Then click the search button:



The Search Result(s) panel displays matching resources, if found.



- a. In the Search Result(s) panel, select the resource you want to add and click **Add** to display it in the Selected Resources panel.
 - b. Click **OK** to add the resource as an attribute.
2. If you know the resource ID, click:

 - a. Enter the resource ID in the text field and press [Enter]. The Search Results panel displays matching resources, if found.
 - b. In the Search Result panel, select the resource you want to add and click **Add** to display it in the Selected Resources panel.
 - c. Click **OK** to add the resource as an attribute in the Edit panel.

Keyboard Shortcuts (Hot Keys)

You can accomplish many actions in the by using the default keyboard shortcuts or *hot keys*, instead of menus and mouse navigation. The standard keyboard shortcuts and their associated actions is listed in the table below.



Tip: You can view the default keyboard shortcut schemas and set up custom shortcuts on the Hot Key tab in the Console Preferences dialog (Console menu option **Edit > Preferences**, click **Manage Hot Keys**). For information on how to view or configure Console keyboard shortcuts, see ["Managing Hot Keys" on page 58](#).

Keyboard Shortcuts

Task	Keyboard Shortcut	Description
Annotate events	Ctrl-T	Select one or more events in any grid view, and use Ctrl-T keyboard command (as an alternative to the right-click Annotate Events menu option). See "Annotating an Event" on page 296 .
Mark events reviewed	Ctrl-R	Select one or more events in any grid view, and use Ctrl-R keyboard command (as an alternative to right-click Mark as reviewed menu option). See "Collaborating on Events (Event Annotation)" on page 295 .
Copy	Ctrl-C	See "Using the Edit Menu" on page 82
Cut	Ctrl-X	See "Using the Edit Menu" on page 82
Delete	Delete key	See "Using the Edit Menu" on page 82
Find	Ctrl-F	See "Using the Edit Menu" on page 82
Open the Edit menu	Alt-E	See "Using the Edit Menu" on page 82
Paste	Ctrl-V	See "Using the Edit Menu" on page 82
Redo	Ctrl-Y	Re-do any text edit operation.
Select All	Ctrl-A	See "Using the Edit Menu" on page 82
Undo	Ctrl-Z	Undo any text edit operation.
Exit/shut down the Console	Alt-F4	See "Using the File Menu" on page 81
Open the File menu	Alt-F	See "Using the File Menu" on page 81
Open the View menu	Alt-V	See "Using the View Menu" on page 82
Open the Window menu	Alt-W	See "Using the Window Menu" on page 84
Open the Tools menu	Alt-T	See "Using the Tools Menu" on page 84

Keyboard Shortcuts, continued

Task	Keyboard Shortcut	Description
Open the System menu	Alt-S	See "Using the System Menu" on page 85
Open the Help menu	Alt-H	See "Using the Help Menu" on page 85
Open the Help directly	F1	See "Using the Help Menu" on page 85

Creating Shortcuts for Resources

For most resources, the Navigator panel contains a **Shortcuts** option above the resource tree. If you click this option, the top panel expands to display Recents and Favorites.

- Recents is a container for the most recent resources you viewed on the Edit panel. ESM automatically populates this list.
- Favorites is a container of your own list of frequently-used resources . You maintain this list.

The Recents and Favorites features are not available for:

- Saved Searches and Search Filters, which are maintained and used in the ArcSight Command Center
- Notifications, which are not resources but groupings of user destinations for notifications

Showing Recently Viewed Resources

The Recents list is automatically populated and updated as you open a resource's Edit panel. The list displays the last 10 you viewed. This list is persisted through all your Console sessions. You can manually remove an item from the list.

To use the Recents options:

1. Click **Shortcut**.
2. Right-click a Recents entry and select an option:

Option	Description
Edit [Resource]	The resource's Edit panel opens.
Delete [Resource]	A prompt appears, confirming the deletion. If you confirm, the resource is deleted from the resource navigator.
Add to Package	The Package selector appears. See "Adding Resources from the Resource Navigator" on page 735 for more details.
Find in Resource Tree	Console expands the resource group in the resource tree Navigator and highlights the resource.
Remove from Recents	The selected resource is removed from the Recents list.

Adding Resources to the Favorites List

Maintaining a Favorites list saves you from expanding nested resource groups to access resources that you use frequently.

You can add up to 10 resources to the Favorites list, one resource at a time. The resources can belong to any resource group as long as they are of the same resource type. For example, 10 active channels you add to the Favorites list can come from any active channel group. This list is persisted through all your Console sessions.

1. On the Navigator's resource tree, click **Shortcuts**.
2. Right-click the resource and select **Add to Favorites**.

To use the Favorites options:

1. Click **Shortcuts**.
2. Right-click a Favorites entry and select an option:

Option	Description
Edit [resource name]	The resource's Edit panel opens.
Delete [resource name]	A prompt appears, confirming the deletion. If you confirm, the resource is deleted from the resource navigator.
Add to Package	The Package selector appears. See "Adding Resources from the Resource Navigator" on page 735 for more details.
Find in Resource Tree	Console expands the resource group in the resource tree Navigator and highlights the resource.
Remove from Favorites	The selected resource is removed from the Favorites list.

Printing from the Console

You can print Navigator trees for all resources. You can print resource definitions for rules, filters, and cases, as well as conditions from the "[Common Conditions Editor \(CCE\)](#)" on [page 887](#) (for all resources with filters). You can print from all grid or channel views.



Tip: You have the option to display a Print Preview dialog before you send your job to the printer. Enable the Print Preview dialog through the Console's Preferences > Global Options menu. See "[Changing Global Options](#)" on [page 49](#) for details.

Printing Navigation Tree Views of Resources

To print the Navigation tree for a resource:

1. In the Navigator, choose the resource you want to print.
2. Click items in the tree to expand or collapse folders in the tree depending on what you want to see in the printout.



Tip: A printout of the Navigation tree for a resource will show the tree exactly as it is displayed on the Console. Folders that are expanded or collapsed on the Console will show the same way in the printout. To print the tree showing the items contained in a particular folder, expand the folder in the Navigation tree before selecting the Print option.

3. Right-click any element in the Navigation tree for that resource and choose **Print <ResourceName> Tree**. (For example, Print Rule Tree.) Regardless of which item you select to access the right-click menu, the whole tree prints.
4. The system displays a print preview that matches the resource's tree view on the Navigator panel.
5. Click **Print** to bring up a standard Print dialog, and set these properties (destination printer, page layout to use, and so on).
6. Click **OK** to print.

Printing Resource Definitions

You can print resource definitions for rules, filters, and cases. You can print a resource definition from the Navigator tree or from within the resource editor.

To print a resource definition:

1. In the Navigator, choose the type of resource you want to print.
2. Right-click an instance of that resource (a rule, filter, or case), and choose **Print <ResourceName> Definition** (for example, Print Rule Definition).

Or

Double-click a resource to open its editor in the Inspect/Edit panel, then right-click the topmost tab in the editor and choose **Print <ResourceName> Definition**.

The system displays a print preview such as the preview of a Rules definition.



Tip: From the Print Preview of a resource definition, you can export the displayed information into an HTML file or send the preview directly to a printer.

3. Click **Print** to bring up a standard Print dialog, and set these properties (which printer, page layout, and so on).
4. Click **OK** to print.

To save the print preview as HTML:

1. On the Print Preview dialog, click the **Export to** tool button.
2. In the file browser, navigate to the location where you want to save the HTML file.
3. Enter a name for the file in the File Name field. The File Type is Web Page (*.html) by default.
4. Click **Save**.

Printing Grid Views

Active channels and active lists are examples of grid views.

To print items from a grid view:

1. Select one or more items in the grid. To select multiple, adjacent items, use the **Shift** key and mouse click, or click and drag. To select non-adjacent items, use the **Alt** key in combination with mouse clicks.
2. Right-click and choose **Print Selected Rows**.

The system displays a preview of the printout.



Note: The format of a grid view printout is determined by the number of columns in the table and the configuration of the Column Flip Limit, which is set in the Console Preferences dialog. For more information, see ["Using Column Flip Limit to Format Grid View Printouts "](#) below.


3. Click **Print** to open the Print dialog, and set these properties (which printer, page layout, and so on).
4. Click **OK** to print.

Printing Conditions Tree Summary

You can print Conditions for any resource with filters.

1. Open the resource in the Editor.
2. Click the **Conditions** tab.
3. Right-click anywhere on the Edit tab in the [Common Conditions Editor \(CCE\)](#).
4. Select **Print Conditions Tree and Summary** from the context menu.

The system displays a preview of the printout. For example, here is a Print Preview of the filter for a rule.




5. Click **Print** to bring up a standard Print dialog, and set these properties (destination printer, page layout to use, and so on).
6. Click **OK** to print.


Using Column Flip Limit to Format Grid View Printouts

For printing tables from Grid Views (channels, lists, and so forth), you can configure the **Column Flip Limit** in the Console Preferences.


Choose **Edit > Preferences**, and click **Grid View Options**. The default is 10 columns.



Grid views with the same or fewer columns than the Column Flip Limit print as a table, the same as that shown in the UI on the Console grid view.



Grid views with more columns than the Column Flip Limit print details per-row rather in a normal table like that shown on the Console grid view.



Instructions for setting the Column Flip Limit for grid views is also summarized in ["Setting Grid Options for the Viewer Panel" on page 52](#).

Saving and Sending Settings

Purpose: The **File Save** and **Save As** options allow you to save Console settings (.ast files) locally. You can also save and load your own personal settings from the ArcSight Manager by using the **File Save to Manager** and **File Load from Manager** options. That way, for example, you can quickly restore Console settings when you move to an Console running on a different computer.

Where: File menu

To save your settings to a file:

The Console saves your settings in the file you specified, on the local computer. Later, you can restore those settings to return the Console to that configuration, using the **File>Open** command.

1. Choose **File > Save** or **File > Save As**.
2. In the Save dialog box, navigate to a directory and enter a file name.
3. Click **Save**.

To save a file to the ArcSight Manager:

Choose **File > Save to Manager**.

Your Console settings are saved to a file (based on your login user name with the .ast extension) and maintained by the ArcSight Manager.

To reload a file from the ArcSight Manager:

On the **File** menu, choose **Load from Manager**. The Console loads the saved settings (.ast) file and asks whether you want to apply them to your current session. If you say **Yes**, the Console restarts and refreshes the display.

To send a file by e-mail:

1. Choose **File > Send To**.
2. In the Send To dialog box, enter the **E-mail Address** and click **OK**.

Error and Warning Messages

Certain error messages, warnings, and notifications appear in a small dialog. To capture the error message and supporting data, click the Copy button or check **Copy message to system clipboard** to copy the entire message to the Clipboard. You can then paste the error message in text fields in the , into the body of an e-mail message, or other applications.

Chapter 3: Managing Users and Groups

Managing User Groups

ESM user groups are designed to contain users with a common set of roles (see the topic, "User Roles," in *ESM 101*) and permissions. ESM provides the following user groups:

ESM User Group Types

Group	Description
Administrators	Associated with the administrator role with all privileges and permissions, including changing other groups' privileges and permissions
Custom User Groups	Minimum privileges and permissions, but administrator can modify.
Default User Groups	Further subdivided into subgroups that map to roles in the enterprise's security operations center (SOC). Each subgroup has a predefined set of privileges but administrator can modify. • Analyzer Administrators. Associated with the author role. Responsible for creating ESM content. • Operators. Associated with the operator role. Use content created by authors to monitor security-related activities. Handle and resolve cases as assigned. • Operators/Analyst. Associated with the business user role.



Tip: If you belong to the Administrators group, you can view all groups and their associated permissions. Right-click a group and choose Edit ACL to open the ACL Editor for that group. Refer to ["Managing Permissions" on page 107](#).

Where: Navigator > Resources > Users

To create user groups:



Caution: Do not exceed more than 10,000 resources in a group.

1. Right-click a group and select **New Group**.
A name text field appears under the group you selected.
2. In the name text field, enter a name.
3. Press **Enter**.
4. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).

To rename user groups:

1. Right-click a group and choose **Rename**.
2. In the name text field, rename the group.
3. Press **Enter**.

To edit user groups:

1. Right-click a group and choose **Edit Group**.
2. In the **Group Editor**, edit the **Name** and **Description** text fields.
3. Click **OK**.

To move or link (copy) user groups:



Note: To copy multiple resources at once, use Copy and Paste. You can drag and drop only one resource at a time.

1. Navigate to a group and drag and drop it into another group.
2. Choose **Move** to move the group or **Link** to create a copy of the group that is linked to the original group.

If you choose **Link**, you create a copy of the group that is linked to the original group. Therefore, if you edit a linked group, whether it is the original or the copy, all links are edited as well. When deleting linked groups, you can either delete the selected group or all linked groups.

To delete user groups:

If you delete a group, the users within that group are also deleted, unless they are also contained by other groups.

1. Right-click a group and choose **Delete Group**.
2. In the dialog box, click **Yes**.

To set Console startup views:

You can define the set of active channel and dashboard resource groups that members of a given ArcSight user group see by default when they first log in. This includes both ArcSight Console and Command Center users. These channels and dashboards are initial defaults only: when users begin changing the content of the Viewer panel, the ArcSight Console and Command Center follow their normal behavior of remembering the most recent state.

The default active channels and dashboards you select for user groups are listed in the User Group Editor on the Startup Views tab.

1. Right-click a user group and choose **Edit Group**.
2. In the User Group Editor, click the **Startup Views** tab, then the **Active Channels** or **Dashboards** tabs.
3. In either resource tab, click **Add** to open a resource selector dialog box.
4. Navigate to and select the appropriate active channels or dashboards to set as users' start-up resources, and click **OK**. Repeat this step to add more resources.
5. Click **Refresh** to update the current list of resources, or click **Remove** to take a selected resource off the list. Click **Edit** to change a selected resource in its own editor.
6. Click **Apply** to make changes and leave the editor open, or click **OK** to apply your changes and close the editor.

The following topics include configuration instructions related to user groups:

- ["Managing Permissions" on page 107](#)
- ["Managing Notifications" on page 216](#)
- ["Managing Notification Destinations" on page 219](#)
- ["Creating or Editing a User" on the next page](#)

Managing Users

You manage numbers of users by organizing them into groups based on roles or other logical groupings, setting their permissions and passwords, and enabling or disabling their login functionality. Permissions to access specific resources (for example, to create rules or reports) are granted to specific groups by editing the access control lists (ACLs) for those groups.

When users log in, they are allowed to perform any operations for which they are granted permission through their membership in one or more groups.

When you create an ESM user, that person automatically receives access to a set of resource groups. Users can store, create, edit, or delete resources within their groups without jeopardizing other users' resources.



Note: Some system operations, for example, audit event generations, are done on behalf of a special system user called 1ROOTUSER. When you are investigating event details, you might see a user ID with this value. This user ID is valid and intended for internal use only.

See the following topics:

- ["Creating or Editing a User" below](#)
- ["Resetting User Passwords" on page 104](#)
- ["Moving or Linking a User" on page 105](#)
- ["Deactivating and Reactivating a User" on page 105](#)
- ["Deleting a User" on page 106](#)

Creating or Editing a User

Where: Navigator > Resources > Users

Procedure

1. Locate the user group appropriate for the user's role (see ["Managing User Groups" on page 99](#) for the group types) .
2. If you are creating a user, right-click the group for this user and choose **New User**.
If you are editing a user, expand the group, right-click the user, choose **Edit User**.
3. In the User Editor's **Attributes** tab, set these fields in the **Login** section:

Login Attributes

Fields	Description
User ID	User name for login ID. This is a required field.
User Type	Choose a user type from the drop-down menu. This is a required field. The currently supported user types are: • Normal User: Has full privileges to use both the ArcSight Console and Command Center, and all tools. Only apply this user type to accounts that actually need access to the ArcSight Manager. • Management Tool: Has only the privileges needed to run certain management tools used in conjunction with network management products. • Forwarding Connector: Has only the privileges needed by the ForwardingConnector. • Archive Utility: Has only the privileges needed to run the archive utility. Access to specific resources is controlled through ACLs. • Connector Installer: A specialized identity used only to add SmartConnectors to the system. • Web User: Has privileges to use the ArcSight Command Center but not the ArcSight Console or other tools.

Login Attributes, continued

Fields	Description
Login Enabled	- Select the Login Enabled checkbox to give the user login privileges (a checkmark indicates this feature is on). - Or leave it deselected and off (no checkmark showing) to disable logins for this user: Note: A user account login must be <i>enabled</i> to allow login access to the Console.
External User ID	Optionally, provide an alternate, external user ID. (An external user ID might be relevant if you have user accounts from other applications feeding into user database.)
Password	Enter a password for this user. This is a required field. By default, passwords require a minimum of 6 characters, can contain a maximum of 20 characters, and can contain numbers, letters, or a combination. System administrators can set special policies or requirements for their sites through a configuration file. For information on password restrictions see the ESM Administrator's Guide . You can modify passwords later. See " Resetting User Passwords " on the next page.
Confirm	Re-type the password to confirm it. This is a required field.

- Set these fields on the Attributes tab in the **User** section:

User Attributes

Fields	Description
Last Name	User's last name
First Name	User's first name
Title	User's job title
Department	User's department
Phone	User's phone number
Fax	User's fax number
E-mail	User's e-mail address. Use the format user@host.domain. The "@" sign and host domain are required. E-mail addresses are not case-sensitive.
Pager	User's pager number, if applicable.



Note: For phone, fax, or pager numbers, parentheses (), dashes (-), and periods (.) are allowed. Alphabetic characters are not allowed.

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see "[Common Resource Attribute Fields](#)" on page 725.

5. Optional: If you created commands to integrate with other applications, set the Integration Parameters attributes for the user if applicable. Refer to ["Integration Commands" on page 664](#) and ["Setting User Login Parameters" on page 683](#) for instructions.
6. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).

A user can belong to different user groups simultaneously. Those groups can have different sets of access rights. Users inherit privileges from their parent groups, so plan your group assignment carefully.

To assign this same user to another group, see ["Moving or Linking a User" on the next page](#).

Resetting User Passwords

Purpose:

- Administrators reset user passwords directly if a user's original password has been compromised.
- Users explicitly request they reset their own password.

Permission required: Administrator privileges

Where: Navigator > Resources > Users > *user group* > *user*

To reset passwords:

- To let users reset their own password, right-click the user and select **Reset Password**.
The ArcSight Manager assigns a new random password (8 characters, including numbers and letters) and sends it to the selected user's assigned e-mail address.
Or
- To reset the user's password yourself, right-click the user and select **Edit User**. In the user edit panel, enter the new password and re-enter to confirm.



Caution: Sending a password by e-mail may not be secure.

Alternatively, use the following command on ArcSight Manager to reset a user's password:

```
arcsight resetpwd
```

Read the following topics in the [ESM Administrator's Guide](#):

- Managing Password Configuration
- ArcSight Commands for more details on the resetpwd command

Moving or Linking a User

A user can belong to different user groups simultaneously. Those groups can have different sets of access rights. Users inherit privileges from their parent groups, so plan your group assignment carefully.

After the user is created, you can assign that user to multiple groups by linking.

Where: Navigator > Resources > Users

1. Navigate to a user and drag and drop it into another group.
2. Choose **Move** to move the user or **Link** to create a copy of the user that is linked to the original user.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

Deactivating and Reactivating a User

A user is deactivated for the following reasons:

- The ArcSight administrator manually clears the user's `Login Enabled` checkbox. The administrator can also right-click and choose **Delete User**, then click the **Disable Login** button instead of **Delete**.
- The user attempted to log in and failed three times.
By default, the user can log in again if 10 minutes have passed after the previous login failure. However, you can set the `auth.auto.reenable.time` to -1 so that after reaching the failure limit, an administrator must manually re-enable the user.
See "Restricting the Number of Failed Log Ins" in the [ESM Administrator's Guide](#).
- The user has been inactive for 90 days.

If you want to change the inactive period, set the `auth.user.account.age` property. Refer "Disabling Inactive User Accounts" in the [ESM Administrator's Guide](#).



Note: If you are changing any property, see "Changing Properties File Settings" in the [ESM Administrator's Guide](#).

A deactivated user is denied access to ArcSight Console and ArcSight Command Center. On the Console, the icon associated with a deactivated user appears gray and the **Login Enabled** checkbox is cleared.



Tip: When a user's login is enabled or disabled, the audit event, **User updated**, is generated. However, the event does not indicate what type of update has occurred. The **Enabled** or **Disabled** information is stored in **DeviceCustomString6** field. One way you can view this information is to add **DeviceCustomString6** as a column to the **System Events Last Hour** channel.

Where: Navigator > Resources > Users

To reactivate a user:

1. Navigate to the deactivated user.
2. Right-click and select **Edit User**.
3. Click **Login Enabled**.
4. Click **Apply**.

The user icon's color is restored.

Deleting a User

Permission required:

- Administrator privileges
- Non-administrators with permission to delete users from within their own group. Refer to ["Granting or Removing User Group Permissions" on page 110](#) for details.

Where: Navigator > Resources > Users

1. Right-click the user and select **Delete User**.



Caution: A dialog confirms if you want to delete or disable the user's login (deactive the user but keep the user definition in the database). Deleting a user means deleting all resources that user created: the user's rules, lists, cases, and so on. Click **More Information** for a list of these resources. If you need these resources, copy them to another resource group before deleting the user. If you need more time, disable the login first to prevent unauthorized access.

2. In the dialog box, click **Delete** to delete the user and the listed user's resources or click **Disable Login** to disable the user.

Chapter 4: Managing Permissions

The tasks of managing users is largely that of managing their access to and use of resources.

Editing Access Control Lists (ACLs)

The user groups ACL Editor has these tabs for viewing or editing permissions on resources, operations, user groups, events, and sortable field sets:

- **Resources** tab - Lists all resources available to the user group with either inspect or edit permissions; lets you add custom resources and edit permissions to those resources.
- **Operations** tab - Lists operations for which this user group has permissions and lets you add and edit **operations** permissions. For example, a user group can have permissions to enable or disable data monitors.
- **User Groups** tab - Lists the user groups with either inspect or edit access to the selected user group; and lets you add user groups.
- **Events** tab - Lists event filters for which this group has permissions and lets you add or remove event filter permissions. This user group is permitted to see and annotate only events from the filters listed in the Events tab. By default, custom user groups inherit their ACL settings for events from the parent group. If the user group has no access to event filters, the behavior is as if the group's specified filter in the ACL editor were *Filters/Shared/All Filters/ArcSight System/Core/No Events*.



Note: To view event data in reports and query viewers, a non-administrator user must have Read access to the /All Filters/ArcSight System group. This permission can be set in the **Resource** tab of the **ACL Editor**. For more information, see ["Granting or Removing Resource Permissions" on the next page](#).

- **Sortable Field Sets** tab - Lists sortable field sets for which this user group has permissions.

See also ["Access Control Lists" on page 810](#).



Caution: Always remember to have both ArcSight Console and ArcSight Command Center users log out and log back in after you change user or resource access permissions, so they can experience those changes.



Tip: The Resource ACL display shows relationships between users and groups, and how permissions are acquired for each of the user groups. Child groups inherit permissions from parent groups. For example, consider the following scenario.

- A user logged in as Administrator (belonging to the group /All Users/Administrators) has read and write permissions by virtue of being in the Administrators group.
- All users have read permissions because they belong to the group /All Users/Default User Groups by default.
- A user logged in as an Analyzer Administrator has both read and write permissions because they inherit read permissions from the parent group (/All Users/Default User Groups) and get write permissions per the Analyzer Administrators child group.

Granting or Removing Resource Permissions



Caution: Be sure to set permissions on resources and permissions on events appropriately for user groups.

Preventing users from viewing groups of resources does not necessarily prevent those same users from viewing event data on those resources.

Users with permissions to view certain events (determined by event filters as described here), can view *all event fields* for those particular events (in reports, query viewers, and so forth) even if they do not have permissions on some *resources* reflected in the event data.

For example, a user with no read permissions on an asset could still have permissions to view event data related to the asset, and thereby have access to the data contained in the event fields (such as server name, IP address) in the context of that event.

As a best practice, keep the above in mind when granting permissions on events. Otherwise, you might give some users a view into resource information through event data that you did not intend for them to see.

Where: Navigator > Resources > Users > *user group*

To grant permissions to resources:


1. Right-click a user group and select **Edit Access Control**.
2. In the **ACL Editor**, select the **Resources** tab.

Available resources are listed based on *user permissions*, therefore the list of resources will not be the same for each group.

3. Add or remove permissions on a resource for this user group as follows.
 - **To edit permissions on a resource *shown* in the current list**, click the (R) read or (W) write checkbox next to a target resource to add or remove permissions on that resource.

A checkmark means that this user group has access to the associated resource. A blank checkbox means this group does not have access to the resource.

- To add permissions for a resource **not shown** in the current list, select a resource from the Resource drop-down menu at the top of the Resources tab and click **Add**.



The resource selector dialog for the chosen resource is displayed. Select the resources you want to add permissions for and click **OK**.

The resource you added is listed as a target on the Resources tab and then you can edit its **Read/Write** permissions as needed.

- To remove a resource from the list (and **remove all permissions on it** for this group), select the resource in the list and click **Delete**. (The Delete button is at the bottom of the Resources tab).

4. Click **OK** on the User Group ACL Editor to save changes to Resources permissions.

Granting or Removing Operations Permissions

Examples of operations are deleting cases, reading and writing fieldsets, and deploying data monitors, among others. Default user groups available in ESM have their own set of permissions to which they have permission, such as case deletion, and so on. Other operations, such as data monitor deployment, require explicit granting of permissions to user groups. See also ["Controlling Who Has Permissions to Deploy Data Monitors" on page 116](#).

Any new groups added under Custom User Group may not have any access to most operations. Administrators can allow or block users for operations permissions by setting permissions on a particular operation.

Where: Navigator > Resources > Users

1. Choose a group.
2. Right-click the user group and choose **Edit Access Control**.
3. In the **ACL Editor**, select the **Operations** tab.

The operations for which this user group has permissions (if any) are listed.

4. Add or remove user group permissions to perform an operation as follows.

- **To add permissions to perform an operation not listed**, click **Add**.
In the Permissions Selector dialog, select the operations (expand the nodes as required) you want to add permissions for and click **OK**.

The list of Operations is updated to include the one you added. Operations listed are those this user group has permissions to perform.

- **To remove permissions to perform an operation**, select the operation in the list and click **Delete**. The Delete button is at the bottom of the Operations tab.

5. Click **OK** on the User Group ACL Editor to save changes to Operations permissions.

Granting or Removing User Group Permissions

Where: **Navigator > Resources > Users > user group**

To grant permission to edit user groups:

1. Right-click the user group and select **Edit Access Control**.
2. In the **ACL Editor**, choose the **User Groups** tab.

The User Groups tab lists all user groups for which members of the selected group have inspect (**Read**) or edit (**Write**) permissions, and lets you add/edit group permissions.

Set Permissions On...		
Resources	Operations	User Groups
Events Sortable Field Sets		
User Group	1	
/All Users/Administrators		




Tip: This is where you grant or deny members of the group you are editing permissions to edit their own user groups. Depending on your own user permissions, some user groups may or may not be shown, and Read/Write checkbox options may or may not be editable.


3. Add or remove permissions on a user group as follows.
 - **To edit permissions on a user group shown in the current list**, click the (**R**) read or (**W**) write checkbox next to a target resource to add or remove edit permissions on that user group.
A checkmark means that this user group can edit permissions on the associated group. A blank checkbox means this group does not have edit permissions on it.

User Group	R	W
/All Users/Administrators	☒	☒


- **To add permissions on a user group not shown in the current list**, click **Add**.



The resource selector dialog for the chosen resource is displayed. Select the groups you want to add permissions for and click **OK**.



The user group you added is now listed on the User Groups tab and then you can edit its **Read/Write** permissions as needed.




- To remove a user group from the list (and **remove all edit permissions on it**), select the user group in the list and click **Delete**. (The Delete button is at the bottom of the User Groups tab).

4. Click **OK** on the User Group ACL Editor to save changes to User Group permissions.

To grant non-administrators permission to delete users:

By default, only administrators have permissions to delete users in a group. If you want to grant non-Administrator users permission to delete users within their group (Group1 used in this example), first provide *Write* access to the group by editing access to **User Groups** in the ACL Editor, as described in the previous procedure to grant permission to edit user groups.

After following the instructions, verify Group1's ACL Editor in the User Groups tab. Group1 should appear on the list, as shown:



Additional settings are required. One of them is setting a server property. The other setting is providing Write access to user Reports. This is because deleting users will also delete the resources they created, including query viewers, reports, and so on. Reports created by that user cannot be deleted unless delete permission for that user's reports is also granted. The following steps provide instructions on the additional settings.

1. Read thoroughly "Managing and Changing Properties File Settings" in the [ESM Administrator's Guide](#). In the `server.properties` file, set the following property:

```
user.allowmodification=true
```

2. Restart the Manager.
3. Log into the ArcSight Console as Administrator, and select the **Users** resource in the Navigator.
4. Select the group for non-administrators (Group1 as an example) who will be allowed to delete users in its own group.
5. Right-click Group1 and choose **Edit Access Control** to display the ACL Editor.
6. On the ACL Editor, click the **Resources** tab.
7. Select **Report** in the Resource drop-down menu, and click **Add** to display the Reports Selector popup.
8. In the Selector popup, select all users under Reports/Shared/Personal/ and select each user belonging to Group1. Click **OK**. All users are shown as Resources targets.
9. Click to set **Read (R)** and **Write (W)** permissions as desired.
10. Click **Apply** or **OK** to save your changes.

Members of Group1, even if they are not administrators, can now log into the ArcSight Console and delete users in their own group. To delete users, refer to ["Deleting a User" on page 106](#).

Adding or Removing Enforced Filters

About:

Enforced filters define which events a user group can view. By default, all new groups cannot view any events. If you view the Events tab on a new group's ACL Editor, the filter is shown as

```
/All Filters/Arcsight System/Core/No Events
```

After you add filters to this tab, these filters become the user group's enforced filters that are enforced at run time. The filters you add can be ArcSight-provided filters or filters you created, based on individual groups' requirements.

By default, members of the administrators group can view all events, as indicated by the Administrators group's enforced filter: /All Filters/ArcSight System/Core/All Events.

Prerequisite:


Event filters must be available before you can add them to the Events tab of the ACL Editor for the user group. For more information about filters in general, see ["Filtering Events" on page 306](#). For more information about events, see ["Events" on page 995](#) and ["Event Categorization" on page 997](#).

Important notes about enforced filters:

- ESM evaluates the enforced filters with an OR operator. Evaluating events against filters using OR becomes relevant especially if different filters are applied to a hierarchy of user groups, or if a user is linked to multiple user groups. You should keep these relationships in mind, to determine the ultimate set of events that a user sees.
- An event only needs to match one of the filters, for that event to be accessible. This means if the ACL Editor's Events tab has multiple filters, not all filters are necessarily evaluated after the first match is found. You can consider combining multiple filters into one filter using the MatchesFilter operator, then add that filter to the ACL Editor's Events tab for the user group, to ensure that all filters are evaluated.
- Active channels, when launched, use the enforced filters associated with the user who launched the channels.
- Reports and query viewers use the enforced filters to return and display data.
- Trends and data monitors use the enforced filters of the user who created these resources.
- Users have the ability to annotate events that match any one of their enforced filters.

Where: Navigator > Resources > Users > *user group*

1. Right-click the user group and select **Edit Access Control**.



2. In the **ACL Editor**, select the **Events** tab.

The default enforced filter is listed on the tab.



Caution: Be sure to set permissions on resources and permissions on events appropriately for user groups.

Preventing users from viewing groups of resources does not necessarily prevent those same users from viewing event data on those resources.


Users with permissions to view certain events (determined by event filters as described here), can view *all event fields* for those particular events (in reports, query viewers, and so forth) even if they do not have permissions on some *resources* reflected in the event data.

For example, a user with no read permissions on an asset could still have permissions to view event data related to the asset, and thereby have access to the data contained in the event fields (such as server name, IP address) in the context of that event.


As a best practice, keep the above in mind when granting permissions on events. Otherwise, you might give some users a view into resource information through event data that you did not intend for them to see.

3. Add or remove user group permissions to view events as follows.


- **To add permissions to view events** captured by a filter not shown in the current list, click **Add**.



On the Filters Selector dialog, choose the filters for events that the user group can view and click **OK**. For example:



The list of enforced filters for the user group is updated to include the ones you added:



The default /No Events filter is disabled as you add enforced filters to the Events tab.

- **To remove enforced filters** (event filters for this user group), select a filter in the list and click **Delete**. The Delete button is at the bottom of the Events tab. You cannot delete the default /No Events filter.

4. Click **OK** on the User Group ACL Editor to save changes to Operations permissions.

Permissions for Sortable Field Sets

ArcSight-provided sortable field sets are used to manage processes that come from all resources. To minimize the impact on performance, you are provided two pre-indexed field sets on which you can sort:

- All Field Sets/ArcSight System/Sortable Field Sets/ Field Set Based on ARC_E_ET Index
- All Field Sets/ArcSight System/Sortable Field Sets/ Field Set Based on ARC_

E_MRT Index



These field sets are indexed for the event's end time (ET) and Manager's receipt time (MRT), respectively. For additional information, see ["Sortable Field Sets" on page 1054](#).

Sharing Resources

Purpose: To share your resources with other users by moving, copying, or linking your resource to or into another resource's Public group.

Where: **Navigator > Resources > *resource* or *resource group***

To share a resource:

1. Drag a resource, for example, a filter you created, or a resource group, and drop it into the Public group.
2. Select one:

Copy	Create a separate copy of the resource that is not affected when the original resource is edited
Link	Create a copy of the resource that is linked to the original resource. Therefore, if you edit a linked resource, whether the original or the copy, all links are edited as well. When deleting linked resources, you can either delete the selected resource or all linked resources.

You can also multiple-select resources with the **Shift** key, and drag-and-drop or keyboard copy-and-paste, to move or link them in another group.



Note: To copy multiple resources at once, use Copy and Paste. You can drag and drop only one resource at a time.

Controlling Who Has Permissions to Deploy Data Monitors

Data monitor deployment is controlled through User Access Control Lists (ACLs). Administrators can allow or block users for data monitor deployment permissions.

Depending on the permissions associated with the user group to which they belong, users may or may not have options available on their ArcSight Consoles to **Enable** (*deploy*) or disable (*un-deploy*) data monitors. (See also ["Enabling or Disabling a Data Monitor" on page 278.](#))

Administrators (all users belonging to the Administrators user group) have permissions to deploy and undeploy data monitors.

Administrators can grant permissions to deploy or disable data monitors for other non-Administrator through the Users resource Access Control Lists (ACLs) editor, as described in ["Granting or Removing Operations Permissions" on page 109.](#) As with user permissions for other resources, these are applied at a user group level. As an administrator, you can grant all users in a given group permission to deploy data monitors. After user groups are set up and appropriate permissions are applied to those groups, you can add new users to appropriate groups, and change access permissions for existing users by moving them in or out of various groups. If you want to allow or disallow a particular user the option to deploy data monitors, move the user in or out of a group that has that permission.




Note: About Write and Deploy permissions

Data monitor deployment is an all-or-nothing permission (it applies to all data monitors), while read and write permissions are specific to each data monitor. So, in some cases a user could have read-only access to one data monitor and read-write access to another. To deploy a data monitor, a user needs *both* deployment permissions and write permissions. Users with permissions to deploy data monitors can deploy only those data monitors for which they have write permissions. (Fields in the data monitor editor are grayed out for all users without write permission.)

To configure data monitor deployment permissions:

1. If needed, set up one or more user groups for non-administrator users to whom you want to control permissions to deploy data monitors. For example, at the simplest level you might have a group for analysts and operators who are allowed to deploy data monitors and another for those you want to block from this option.
See ["Creating or Editing a User" on page 102](#) and ["Managing User Groups" on page 99](#) for information on adding, deleting, and editing users and user groups.
2. Follow the instructions provided in ["Granting or Removing Operations Permissions" on page 109](#) to grant or remove permission to deploy data monitors to a particular group. As a part of these instructions, you'll select the **Users** resource in the navigator, right-click a group and choose **Edit Access Control**.
3. In the ACL Editor, click the Operations tab, and click **Add**.
4. On the Permissions Selector, select **Deploy** under Permissions\Shared\All Permissions\ArcSight System\Data Monitor\ and click **OK** to save the settings and close the dialog.

The list of Operations is updated to include deployment permissions on data monitors.



To remove the permission for this group, select the permission and click **Delete**.

5. Click **OK** on the ACL Editor to save your changes.

For information on deploying or disabling data monitors, see ["Enabling or Disabling a Data Monitor" on page 278](#).

For more information on administrator tasks of working with user permissions and ACLs, see ["Managing Permissions" on page 107](#).

How Upgrades Affect Data Monitor Deploy Permissions

Upon installation and deployment of a different version of software (for example, version or patch upgrades), only administrators keep permissions to deploy and disable data monitors. Non-administrators users do not have deploy permissions on data monitors even if they had such permissions as part of the previous configuration.

After upgrades, all users have access to already-deployed data monitors. But, initially, non-administrator users do not have permissions to enable or disable data monitors, nor have access to new data monitors unless an administrator enables (deploys) these.

To re-establish data monitor deployment permissions for non-administrators users after an upgrade, administrators can reconfigure fine-grained permissions. They can re-group users and perhaps link non-administrator users into existing or new groups with more permissions (like data monitor deployment), as described in ["Controlling Who Has Permissions to Deploy Data Monitors" on page 116](#).

Deployment Permissions on Imported Data Monitors


If a user without data monitor deploy permissions imports a data monitor that was archived in the enabled state, the import succeeds but the data monitor is disabled. After the import, the user does not have permissions to deploy the data monitor unless an administrator reconfigures permissions for that user.

If a user with data monitor deploy permissions imports a data monitor that was archived in the enabled state, the import succeeds and the data monitor keeps its enabled (*deployed*) setting. After the import, this user can view the data monitor and re-set its deployment state as needed.

Chapter 5: Modeling the Network

The following topics explain how to model your network and configure various aspects of the network model (assets, locations, zones, and so on), and how to manage customer accounts (if applicable).

ArcSight operates on a data model that enables you to build a business-oriented view of data derived from physical information systems. These distinctions help to clearly identify the events in your network, and provide more layers of detail to correlation. Modeling your network and its assets is part of setup and ongoing maintenance.



The network model consists of the asset model and the network model, which, combined, facilitate building detailed correlation criteria. All of the Network Modeling resources, except Customers, are available as part of the Assets resource.

- The ["The Network Model" below](#) is a representation of the nodes on your network and certain characteristics of the network itself.
- The ["Asset Model" on page 126](#) describes attributes of the assets themselves for different purposes.



Tip: For an overview and information about types of assets and the network model, refer to the "Network Model" section of *ESM 101*.

The Network Model

The network model is a representation of the nodes on your network and certain characteristics of the network itself.

Before you can make an informed decision about what to do about a particular event, it helps to know something about the event's source and destination. Is the source a previous attacker, does it come from a hostile region of the world, or is it a trusted server that has become the

source of an attack? Does the destination host critical applications, or is it a known server of forbidden services?

This kind of information is captured by modeling the assets on your network and particular pertinent attributes of the network. The network model represents information for individual assets and whole zones. For critical assets on the protected network, network modeling captures important facts that help inform your decisions, such as:

- All open ports
- The operating system running on that host
- Known vulnerabilities that might be exposed
- Applications present
- The missions these applications support and their criticality to your operation

For less critical assets, such as a block of addresses on the Internet, it may be sufficient to know general information about them, such as the country in which those assets reside.

The Network Model consists of the following resources. All of these resources, except Customers, are part of the Assets resource.

- **Assets** represent individual nodes on the network, such as servers, routers, and laptops.
- **Asset Ranges** represent a set of network nodes addressable as a contiguous block of IP addresses.
- **Zones** represent portions of the network itself that are characterized by a contiguous block of addresses.
- **Networks** provide an additional distinction to differentiate between two private address spaces with overlapping IP address ranges.
- **Customers** describe the internal or external cost centers or separate business units associated with networks, if applicable to your business environment. Customer tagging is a feature developed mainly to support Managed Security Service Provider (MSSP) environments, although it can also be used by private organizations to denote cost centers, internal groups, or subdivisions. The Customer designation keeps event traffic from multiple cost centers or business units separately identified. Think of a customer as the "owner" of an event, rather than the source or target of an event.

Assets

An asset is any network endpoint with an IP address, MAC address, host name, or external ID. For network modeling purposes, an asset is any endpoint you consider significant enough to characterize with details that make correlation and reporting more meaningful.

Automatically-Created Assets

The system automatically creates assets to model the network nodes that host ArcSight components (Managers, Consoles, and SmartConnectors). It also automatically creates assets for events received from device endpoints on your network that do not already have assets modeled in ArcSight, and, if applicable, for assets arriving from scan reports sent by vulnerability scanners brought in by scanner SmartConnectors. This auto-asset creation feature could require configuration, depending on the assets reporting to the Manager.

Depending on which method you use, assets are placed in the following locations:

- Assets that are created through scanners are placed in the Resource tree under Assets/All Assets/<Zone Group>/<Zone>.
- Assets that are auto-created by any other type of SmartConnectors are placed under Assets/All Assets/ArcSight System Administration/Devices.

As a configuration option, you can also configure it to create assets for devices reporting through SmartConnectors.

Auto-Created Assets for Components

The system automatically creates assets to model the network nodes that host components. These assets do not contain vulnerability information, and are used for system administration.

Component		
ArcSight Manager		An asset for the Manager is added (if needed) every time the Manager service starts.
s		An asset is added for each the first time it connects with the Manager.
SmartConnectors		An asset is created for SmartConnectors only when the SmartConnector begins reporting base events from the device it represents. A Connector can be successfully added to the Manager, but until it starts reporting events from the device it represents, an asset cannot be created for it in the Asset Model. It creates assets differently for SmartConnectors in static zones and those in dynamic zones. For more about static and dynamic zones, see "Zones" on page 124. For details about creating assets for SmartConnectors, see "Creating Assets for SmartConnectors" on page 149.

Devices Discovered by a Vulnerability Scanner

The system also imports asset and vulnerability information from vulnerability scanner reports generated by products such as Nessus, FoundStone, and ISS Internet Scanner. Asset information is passed to the Manager via the scanner SmartConnector appropriate for your vulnerability scanner product based on IP address, MAC address, and host name.

Updated vulnerability information is added to existing assets with matching identifiers. If a matching asset does not already exist, the system creates one.

The system creates assets from vulnerability scan reports differently for dynamic and static zones. For more about dynamic and static zones, see ["Zones" on page 124](#).

For details about how the system creates assets from vulnerability scans, see ["Creating Assets from a Vulnerability Scan Report for Dynamic Zones" on page 149](#).



Tip: Scanner reports list only information received through the scanner, whereas Asset Editors include the full list of both scanner data and vulnerability mappings stored in the system. Therefore, the Editors might show more or different information than the information from scanner reports.

Devices Reporting Through SmartConnectors

The administrator can configure asset creation for each device that reports to that SmartConnector based on IP address, MAC address, and host name when the Manager receives events from SmartConnectors.

This feature makes it possible to add assets to the network model that may not be part of a regular asset scanning report without having to create them individually. Assets created using this method do not contain vulnerability information, although once they are added to the network model, they can be supplemented with matching data that arrives from a scanner report or that you add individually using the .

Administrators can enable the option to create assets for network devices in the Manager Configuration Wizard. See “Running the Manager Configuration Wizard” in the [ESM Administrator's Guide](#).

The system creates assets differently for devices in static zones and those in dynamic zones. For more about static and dynamic zones, see ["Zones" on page 124](#).

For details about how the system creates assets for devices reporting through SmartConnectors, see ["Creating Assets for Network Devices" on page 151](#).

For more about how to tune asset auto creation from the , see the [ArcSight Administration and ArcSight System Standard Content Guide](#). For information about an optional ArcSight Foundation, refer to the Standard Content Guide for that Foundation.

You can customize how the asset auto-creation function works by modifying settings in the `server.properties` file. For additional details, see ["Creating Assets from a Vulnerability Scan Report for Dynamic Zones" on page 149](#). For more about working with properties files, see “Managing and Changing Properties File Settings” in the [ESM Administrator's Guide](#).

For an overview of the ways by which the network model can be populated with assets, see ["Populating the Network Model with Assets" on page 128](#).

Asset Aging and Model Confidence



Note: Only the assets belonging to the following categories are considered for aging:

- /All Asset Categories/Site Asset Categories/Scanned/Open Ports
- /All Asset Categories/Site Asset Categories/Scanned Vulnerabilities

The asset aging function keeps track of the last time an asset was scanned, and incrementally diminishes an asset's model confidence in the priority formula over time to zero if it hasn't been scanned in more than 120 days. (You can configure the time range.)

An asset's age is tracked by default. You can opt to automatically disable an asset that exceeds the configured age limit. This process is described in "Asset Aging" in the [ESM Administrator's Guide](#).



Note: Resolving zone information on disabled assets

To ensure that events get sorted properly, the system continues to resolve an asset's zone information and add it to the event, even when the asset is inactive (disabled).

To see why an asset was disabled:

1. In the Navigator panel, go to the Assets tab in the Assets tree. The disabled asset appears with a grey icon.
2. Right-click the disabled asset and select **Show Disabled Reason**. The message displayed indicates how many days it has been since the asset's last scan.

To re-enable a disabled asset:

If an asset has been automatically disabled, you can manually re-enable it. In the Navigator panel in the Assets tab of the Assets tree, right-click the disabled icon and select **Enable**.

Asset Ranges

An asset range is a group of network assets that use a contiguous block of IP addresses. An asset range is useful if you have many network nodes that would be impractical to track individually, or that may come and go from the network, such as desktop PCs and laptops.

When a SmartConnector processes an event, the Manager, or the correlation engine, its endpoints are identified either as a single asset or as an asset belonging to an asset range. A reference to the asset or asset range identifier is populated in the event schema.


Zones

Zones are ArcSight resources that represent a functional part of the network with contiguous IP addresses, such as DMZ, VPN, wireless LAN, or DHCP.

Every asset or address range is associated with a zone. ArcSight is configured with the standard global IP address ranges represented as zones, so if your network uses only these public IP addresses, ArcSight can resolve them without setting up additional zones.

Zone groups are folders in which one or more zone resource is stored. Although the assets contained in a zone do not inherit the properties of a zone, the zone groups are hierarchical, which means that properties assigned to a zone group apply to all the zones contained within that group.

The following zones are standard:



Create your own zones if you have overlapping private networks. Private networks usually model a functional group within your network or a subnet, such as a wireless LAN, the engineering network, the VPN, or the DMZ.

For details about using the zone editor, see ["Managing Zones" on page 168](#).

Dynamic and Static Zones

The asset auto-creation feature (see ["Asset Auto-Creation" on page 147](#)) relies on zones that are already in place before device discovery occurs, either in customer-created zones, or in default zones. When you add a SmartConnector, you assign one or more existing Networks to that Connector. All assets reported by that Connector are then associated with that Network and the zones the Network represents.

The system differentiates dynamic and static zones to classify the represented asset types.

Static Zones

Devices in a static zone use static (constant) IP addresses. These are devices that stay on the network and use the same IP address for all traffic. In order to identify assets in static zones, the assets must have either a unique IP address, a unique host name, or both.

Dynamic Zones

Devices in a dynamic zone use dynamic addressing (such as DHCP). Dynamic zones represent assets that come and go from the network, such as laptops. By default, the system requires either a MAC address or a host name to identify assets in dynamic zones. The system first looks for a MAC address; if not available, the host name is used.



Caution: Classifying Zones as Static or Dynamic

It is important that zones are classified properly as dynamic or static.

If a zone is classified as static, but hosts assets that come and go from the network, the system may not be able to update the network model properly. For example:

- The updated network might have duplicate and disabled assets.
- Other information, such as vulnerability information and open ports, might not be updated properly.

Static Assets in Dynamic Zones

If an asset is classified as static, but belongs to a dynamic zone, the system treats the asset as if it was in a static zone. See the description and links above for how the asset auto-creation feature works for static zones.

Networks

Networks are ArcSight resources that are used to differentiate between zones whose IP ranges overlap, such as when branch locations assign the same private address spaces to resources used in other corporate locations.

The system comes configured with two standard networks: /All Networks/ArcSight System/Local and /All Networks/ArcSight System/Core/Global. The Local network is where you add your custom zones. Zone mappings in the Local network override the default zone mappings provided by the Global network.

The Global network provides default zone mapping if no local networks are defined, and automatically provides the correct addressing information to ArcSight SmartConnectors when they are installed.

Custom networks are also used to compartmentalize Customer designations in MSSP situations.

When you associate a customer or a location with a network in the Network Editor, zones automatically access this information. (See ["Managing Networks" on page 170.](#))

Asset Model

The resources that make up the asset model are part of the overall network modeling process. The asset model resources describe attributes of the assets themselves for different purposes. *Locations* and *Vulnerabilities* are part of the Assets resource.

Locations

The system provides a location database that maps an IP address to the owning body for the block of IP addresses to which it belongs. Your organization may have finer-grained detail, such as the physical location of all of your networks or networks outside your control, or corrections to the database that the system supplies. The Location resource is the way you can override the default location mappings with location information relevant to your network.

Location is an attribute you can set if the asset you are modeling resides in a geographic location that differs from the location set by the mapping database that associates IP addresses with location information.

Vulnerabilities

The asset vulnerabilities on your network are normally discovered and updated by scanners. The most common manual change to a vulnerability resource is to associate it with a Knowledge Base article. You can associate assets with vulnerabilities from either the Vulnerabilities or Assets editors. (See ["Managing Vulnerabilities" on page 162](#) for information on the vulnerability editor.)

Asset Categories

Asset categories are ArcSight resources that describe the properties of an asset in terms of how it is used. Asset categories are one of the key ways to add differentiation, relevance, and context to the millions of events passing through your network.

Asset categories establish identity, ownership, and criticality of the assets on your network. The root of a particular category (for example, **Criticality** in the group /All Asset Categories/System Asset Categories/Criticality) defines the property itself, whereas the members of the category (for example, the criticality levels Very High, High, and so on) define the possible values for that property.

You can create new asset categories as a right-click option in the navigation panel, and associate categories with assets using the Asset Editor. Most methods for populating the network model described in ["Populating the Network Model with Assets" on the next page](#) include adding asset categories to your assets, asset ranges, asset groups, and zones.

Asset Categories Assigned to Assets, Asset Ranges, and Asset Groups

Categories assigned to individual assets and asset ranges apply only to those individual assets. This is the most granular level to which you can apply asset categories. If an asset falls into an asset range, it inherits the asset categories assigned to the asset range.

Asset Groups are a folder containing one or more Asset resources. Asset Groups are hierarchical, so properties assigned to an Asset Group apply to all the assets in that group.

Categories assigned to asset groups apply to all assets and asset ranges in that group. Individual assets and asset categories within a group inherit the categories assigned to the group, if any, in addition to the asset categories assigned to them individually.

Asset Categories Assigned to Zones

Categories assigned to zones describe the network itself, not assets within it. Use this to categorize traffic on a network where the assets are not constant, such as a wireless or VPN network. For example, categories might describe whether the network is wireless, encrypted, or a VPN. You may be characterizing the network or the traffic on the network (wireless describes the network; encrypted describes the traffic) rather than the assets. Asset categories assigned to zones are not passed on to assets contained within that zone.

For more about asset modeling, see the topic “Asset Model” in *ESM 101*. For instructions about how to set asset categories, see the following topics:

- ["Populating the Network Model with Assets" below](#)
- ["Populating the Network Model Using the Wizard" on page 132](#)
- ["Managing Asset Categories" on page 171.](#)



Caution: Always exercise caution when deleting or changing existing asset categories. Changing an asset category can break existing conditions that use that category. As a best practice, create new categories in new groups.

Populating the Network Model with Assets

There are several ways to populate the network model with the assets that represent your monitored network. Most enterprises use a combination of these methods:

- ["ArcSight Console-Based Methods" on the next page](#)
- ["SmartConnector-" on page 130](#)
- ["ArcSight-Assisted Methods" on page 131](#)



Caution: Do not import assets with an ampersand (&) in the name. The ArcSight resource framework does not support that character in asset and zone names.


When importing assets using a scanner import connector, the automatically-created asset group name is based on the new asset's zone name. If that name already exists in the same folder, then instead of importing the asset there, as you would expect, the connector imports the asset into either a standard System Asset Group or a higher level custom Asset Group to prevent a folder name conflict.

Ensure that if you manually create asset range names, they do not match the Zone names of assets that you intend to import using an asset scanner connector.

ArcSight Console-Based Methods

The provides two ways to populate the network model:

- Manually configuring individual network modeling resources
- Using a Network Modeling Wizard



All the tools for modeling the network are in the . The Network Modeling Wizard provides a quick way to add basic assets to your Network Model at setup time.

Manually, Using Network Modeling Resources

Set parameters for every asset using the network modeling resources (Assets, Asset Ranges, Zones, Networks, and Customers) and asset modeling resources (Asset Categories, Vulnerabilities, and Locations).

Use these tools in conjunction with the other batch-loading methods that only offer limited distinctions. As long as primary identifiers, such as IP address, host name, and MAC address, remain the same, the automatic update methods only update fields with new information so the Network Model remains stable.

See the topic “ArcSight Network Model” in *ESM 101*.

In a Batch Using the Network Modeling Wizard

The provides a Network Modeling wizard as a set-up and configuration tool (menu option **Tools > Network Model**). The Network Modeling wizard enables you to load Assets, Asset Ranges, and Zones along with Asset Category information. If you also add a vulnerability scanner as described in "[SmartConnector-](#)" below, the existing assets in the model are updated with the vulnerability scan report data.

The Network Modeling Wizard is flexible. It can take output from any device type in CSV format. The CSV file can include as many new or pre-existing asset categories as are relevant to the devices without having to add asset category information one by one later using the Asset Category resource in the . This tool is appropriate for initial set-up and configuration, not as a method for maintaining the network model.

For more about the Network Modeling Wizard, see "[Populating the Network Model Using the Wizard](#)" on page 132.

SmartConnector-

These methods enable batch loading and automatic ongoing maintenance. Both methods offer limited distinctions.

 Assets	 Assets
 Vulnerabilities	 Asset Categories
 Asset Categories (OS and open port only)	 Locations
☐ Batch load? Yes	☐ Batch load? Yes
☐ Allow ongoing maintenance? Yes	☐ Allow ongoing maintenance? Yes
 Scanner SmartConnector	 Asset Model Import FlexConnector
 Connector-based methods	

Using the Asset Model Import FlexConnector

The system offers an Asset Model Import FlexConnector that reads Asset, Location, and Asset Category information from a CSV file, which it then sends to the Manager. New assets are added and existing assets in the model are updated .

This method does not create asset ranges, and assumes that Zones, Networks, customers, and locations are already created.

For more about the Asset Import FlexConnector, see the *Developer's Guide for Asset Model Import FlexConnector*.

Automatically From a Vulnerability Scanner Report

Set up a scanner SmartConnector (such as FoundStone, ISS Internet Scanner, or Nessus) to use the output of a vulnerability scan to convert device information into Assets along with Vulnerability information, and basic Asset Categories, such as operating system and open ports. The scanner connector that corresponds with your vulnerability scanning product sets up a directory that to regularly scan for updated reports. It then converts the scanner report output into internal scanner meta-events, which the Manager converts into Assets, open port and OS Asset Categories, and Vulnerabilities. For more about the architecture of how this works, see the topic “How Vulnerability Scans Populate and Update the Network Model” in *ESM 101*.

You can also set the scanner SmartConnector to save network model data as a CSV file, which you can then upload into the Manager using the Files resource during your initial network model setup.


Data derived from vulnerability scanner reports does not create asset ranges, and assumes that Zones and Networks are already created. Once scanner data is imported, you can add Customer and Location distinctions to the assets individually. For details about how the system adds updated vulnerability information arriving from a new scanner report, see ["Reporting on Output from Vulnerability Scanners" on page 168](#).

This method is appropriate for updating and maintaining your network model. Subsequent scans update the basic Asset, Asset Category, and Vulnerability information without overwriting the other network modeling settings you add individually.

For more information about the scanner SmartConnector for your vulnerability scanning product, see the SmartConnector Configuration Guide that corresponds with your vulnerability scanning equipment.

ArcSight-Assisted Methods

ArcSight Professional Services can help you populate the Network Model from an existing configuration database.



As an Archive File From an Existing Configuration Database

Many enterprise networks have third-party systems that already model the properties of the assets on your network. With the help of ArcSight Professional Services, you can export these network models, translate the format into the schema using an ArcSight resource-generating utility, and import it to the Manager as a resource archive with the help of ArcSight Professional Services.

The tools ArcSight Professional Services use can generate any type of resource, so using this method, you can have a fully populated network model without having to do any individual configuration.

Populating the Network Model Using the Wizard

The Network Model Wizard (menu option **Tools > Network Model**) makes it possible to quickly populate the network model by batch loading asset and zone information from Comma Separated Values (CSV) files.



Caution: Before using the Network Model Wizard :

- Make sure you have Administrator privileges.
- Do not import assets with an ampersand (&) in the name. The ArcSight resource framework does not support that character in asset and zone names.

The following data can be imported into the Manager from CSV files:

- **Zones** define functional parts of a network, such as a wireless LAN, an engineering network, a VPN or a DMZ. For the column types of the zones CSV file, see ["Zones CSV File Format" on page 137](#).

- **Assets** are individual nodes on the network, such as servers and routers. For the column types of the assets CSV file, see ["Assets CSV File Format" on page 139](#).
- **Asset ranges** are sets of network nodes addressable as a contiguous block of IP addresses. They are useful when you have many network nodes that are impractical to track individually, or that may come and go from the network, such as laptops. Asset ranges should be a subset of the IP address ranges defined for zones. For the column types of the asset ranges CSV file, see ["Asset Ranges CSV File Format" on page 142](#).

You can import combinations of input CSV files at one time using the Network Model Wizard but only one file of each type can be imported during a single import. For example, if you only have assets to import, you can import only an assets CSV file. If you have a zones CSV file, an assets CSV file, and an asset ranges CSV file to import, you can import all three at once using the Network Model Wizard.

Specifying CSV Column Types

Each CSV file type defines a set of required columns and optional columns. The CSV file can contain columns that are not used by the Network Model Wizard. Columns can be in any order but the Network Model Wizard requires that you specify their types so the wizard knows how to interpret them. Specify the column type using one of the following methods:

- Specify the column type in the header of the CSV file itself, prior to launching the Network Model Wizard. For instructions, see ["Specify the Column Type Using a Header" below](#).
- While running the Network Model Wizard, assign the appropriate column type for each column in the Select Column Headers panel. For instructions, see ["Assign the Column Type in the Wizard" on page 135](#).

Columns not used by the Network Model Wizard must be assigned the column type Ignore. Only columns of type Ignore and Category URI can be repeated in the CSV file. For all other column types, only one instance of the column type can be assigned in the file. If duplicate columns of a non-repeatable column type exist in the CSV file, one of the columns should be assigned the Ignore column type. In a zones CSV file for example, if two name columns appear in the CSV file, assign one to the Name column type and the other to the Ignore column type.

Specify the Column Type Using a Header

In this method, you specify the column type in the first row (header) of the CSV file itself before importing the CSV file using the wizard. The column name in the header must match the column type specified in:

- ["Zones CSV File Format" on page 137](#)
- ["Assets CSV File Format" on page 139](#)

- ["Asset Ranges CSV File Format" on page 142](#)

As shown in following sample zones CSV file, the column names in the first row match the column types specified in ["Zones CSV File Format" on page 137](#). The wizard determines how to interpret each column using the column type specified in the header.

```
Name,Start Address,End Address,Dynamic,Category URI
DMZ Public,<Starting-IP-address>,<Ending-IP-address>,FALSE,
DMZ Corporate,<Starting-IP-address>,<Ending-IP-address>,FALSE,/All Asset
Categories/Site Asset Categories/Business Impact Analysis/Network
Domains/Email/
LAN Corporate,<Starting-IP-address>,<Ending-IP-address>,TRUE,
```

When this zone's CSV file is imported into the wizard, the wizard correctly matches the column types because you specified them in the header.

Select Column Headers for the Zone Data

Click each column heading and select the appropriate column type from the drop-down list. Only the first 15 rows of data are displayed but all the data will be imported. Each column must be assigned a column type.
For a description of each of the column headings, click the Help (?) button.

Name	Start Address	End Address	Dynamic	Category URI
DMZ Public	192.0.2.0	192.0.2.255	FALSE	
DMZ Corporate	198.51.100.0	198.51.100.255	FALSE	/All Asset Categ...
LAN Corporate	203.0.113.0	203.0.113.255	TRUE	

Specifying Multiple Categories in one Category Column

When manually creating a CSV file for importing Zones, assets or Asset Ranges, you can specify more than one category to appear in the same category column:

Column Type - Importing Zones CSV File using Header Row:

```
Name,Start Address,End Address,Dynamic,Category URI,Category URI
HRZone,<Starting-IP-address>,<Ending-IP-address>,FALSE,/All Asset
Categories/ArcSight System Administration/Databases,/All Asset
Categories/ArcSight System Administration/Window Servers
ITZone,<Starting-IP-address>,<Ending-IP-address>,FALSE,/All Asset
Categories/ArcSight System Administration/Databases,/All Asset
Categories/ArcSight System Administration/Window Servers
```

Importing Zones CSV file without Header row is the same, except without the first line.

The above is an example of importing zones, but using multiple categories works the same way for importing assets and asset ranges.

Assign the Column Type in the Wizard

In this method, you assign the column type in the Select Column Headers panels while running the wizard. When the following sample zones CSV file (which does not contain a header row) is imported, the wizard does not know how to interpret all the columns as shown below.

```
DMZ Public,<Starting-IP-address>,<Ending-IP-address>,FALSE,  
DMZ Corporate,<Starting-IP-address>,<Ending-IP-address>,FALSE,/All Asset  
Categories/Site Asset Categories/Business Impact Analysis/Network  
Domains/Email/  
LAN Corporate,<Starting-IP-address>,<Ending-IP-address>,TRUE,
```

Select Column Headers for the Zone Data

Click each column heading and select the appropriate column type from the drop-down list. Only the first 15 rows of data are displayed but all the data will be imported. Each column must be assigned a column type.
For a description of each of the column headings, click the Help (?) button.

Name	Select	Start Address	End Address	Dynamic	Category URI
DMZ Public	This zone de...	192.0.2.0	192.0.2.255	FALSE	
DMZ Corporate	This zone de...	198.51.100.0	198.51.100...	FALSE	/All Asset Ca...
LAN Corporate	This zone de...	203.0.113.0	203.0.113...	TRUE	

By default, when this sample data is imported into the wizard, the second column is automatically assigned to the Select column type but it is a description of the zone and should be assigned the Ignore column type. To change the column type, click the title of the column, and from the drop-down menu select the appropriate column type from the list of options.

Select Column Headers for the Zone Data

Click each column heading and select the appropriate column type from the drop-down list. Only the first 15 rows of data are displayed but all the data will be imported. Each column must be assigned a column type.
For a description of each of the column headings, click the Help (?) button.

Name	Categor...	Start Address	End Address	Dynamic	Category URI
DMZ Public	Ignore	192.0.2.0	192.0.2.255	FALSE	
DMZ Corporate	Name	198.51.100.0	198.51.100...	FALSE	/All Asset Ca...
LAN Corporate	Start Address	203.0.113.0	203.0.113. ...	TRUE	
	End Address				
	Dynamic				
	Category URI				

Zones CSV File Format

Zones define functional parts of a network, such as a wireless LAN, private networks, or subnets. For example, the following network areas could be identified as a zone: the VPN, the DMZ, or an engineering network. Zones are identified with a contiguous block of addresses.



Caution: Each zone should specify a unique range of IP addresses. The IP addresses specified by zones should not overlap. If you import a zone that overlaps with a zone already specified on the ArcSight Manager and the new zone has a different name than the existing zone, the following occurs:

- The new zone is created.
- The existing zone is invalid and is displayed with the broken zone icon in the .

You can define a set of zones in ESM by batch loading zone definitions from a zones CSV file. Zones CSV files contain the columns listed in the table below. When a zones CSV file is selected for import, by default only the first fifteen rows of data are displayed in Select Column Headers for the Zone Data panel. However, when the data is imported into the ArcSight Manager, all the rows are imported. For more information, see ["Increasing the Number of Displayed Rows" on page 143](#).

For the wizard to determine how to process the imported data, the type of each column must be specified. For more information, see ["Specifying CSV Column Types" on page 133](#).

When the Next button is clicked in the Summary of Data to Import panel, the zone data is imported into the ArcSight Manager. The new zones are created in the /All Zones/Site Zones group. For example, if a zone called DMZPublic was specified in the imported zones CSV file, a new zone is created at the following URI: /All Zones/Site Zones/DMZ Public. The new zones are assigned to the default network called Local.

Zone CSV File Columns

Column Type	Description	Required Column?	Repeatable Column?	Example Value
Name	A descriptive name for the zone such as the purpose or geographical location.	Yes	No	DMZ Public
Start Address	The start of the range of IP addresses that defines the zone.	Yes	No	192.0.2.0
End Address	The end of the range of IP addresses that defines the zone.	Yes	No	192.0.2.24

Zone CSV File Columns, continued

Column Type	Description	Required Column?	Repeatable Column?	Example Value
Dynamic	Determines whether the devices defined in the zone use dynamic addressing: • true—devices in the zone use dynamic addressing (DHCP) • false—devices in the zone use static IP addressing	No Default is false	No	false
Category URI	The asset category to assign to zone. NOTE: The wizard does not create new categories. For the category to be assigned, it must already exist.	No	Yes This column can be repeated because a zone can be categorized into more than one asset category.	/All Asset Categories/Site Asset Categories/Business Impact Analysis/Business Role/Service/Web/
Ignore	The column contains data that is not used by the Network Model Wizard when creating zones. For example, this column could contain a description of the zone.	No	Yes	This zone defines the public subnetwork of the DMZ.

An Example of a Zones CSV File

Here is an example of the Zones CSV file:

```
HRZoneA,<Starting-IP-address>,<Ending-IP-address>,FALSE,/All Asset  
Categories/ArcSight System Administration/Databases/
```

```
IT Zone,<Starting-IP-address>,<Ending-IP-address>,TRUE,/All Asset  
Categories/ArcSight System Administration/Databases/
```

Zones CSV File Format

Zones define functional parts of a network, such as a wireless LAN, private networks, or subnets. For example, the following network areas could be identified as a zone: the VPN, the DMZ, or an engineering network. Zones are identified with a contiguous block of addresses.



Caution: Each zone should specify a unique range of IP addresses. The IP addresses specified by zones should not overlap. If you import a zone that overlaps with a zone already specified on the ArcSight Manager and the new zone has a different name than the existing zone, the following occurs:

- The new zone is created.
- The existing zone is invalid and is displayed with the broken zone icon in the .

You can define a set of zones in ESM by batch loading zone definitions from a zones CSV file. Zones CSV files contain the columns listed in the table below. When a zones CSV file is selected for import, by default only the first fifteen rows of data are displayed in Select Column Headers for the Zone Data panel. However, when the data is imported into the ArcSight Manager, all the rows are imported. For more information, see ["Increasing the Number of Displayed Rows" on page 143](#).

For the wizard to determine how to process the imported data, the type of each column must be specified. For more information, see ["Specifying CSV Column Types" on page 133](#).

When the Next button is clicked in the Summary of Data to Import panel, the zone data is imported into the ArcSight Manager. The new zones are created in the /All Zones/Site Zones group. For example, if a zone called DMZPublic was specified in the imported zones CSV file, a new zone is created at the following URI: /All Zones/Site Zones/DMZ Public. The new zones are assigned to the default network called Local.

Zone CSV File Columns

Column Type	Description	Required Column?	Repeatable Column?	Example Value
Name	A descriptive name for the zone such as the purpose or geographical location.	Yes	No	DMZ Public
Start Address	The start of the range of IP addresses that defines the zone.	Yes	No	192.0.2.0
End Address	The end of the range of IP addresses that defines the zone.	Yes	No	192.0.2.24

Zone CSV File Columns, continued

Column Type	Description	Required Column?	Repeatable Column?	Example Value
Dynamic	Determines whether the devices defined in the zone use dynamic addressing: true—devices in the zone use dynamic addressing (DHCP) false—devices in the zone use static IP addressing	No Default is false	No	false
Category URI	The asset category to assign to zone. NOTE: The wizard does not create new categories. For the category to be assigned, it must already exist.	No	Yes This column can be repeated because a zone can be categorized into more than one asset category.	/All Asset Categories/Site Asset Categories/Business Impact Analysis/Business Role/Service/Web/
Ignore	The column contains data that is not used by the Network Model Wizard when creating zones. For example, this column could contain a description of the zone.	No	Yes	This zone defines the public subnetwork of the DMZ.

An Example of a Zones CSV File

Here is an example of the Zones CSV file:

```
HRZoneA,<Starting-IP-address>,<Ending-IP-address>,FALSE,/All Asset  
Categories/ArcSight System Administration/Databases/
```

```
IT Zone,<Starting-IP-address>,<Ending-IP-address>,TRUE,/All Asset  
Categories/ArcSight System Administration/Databases/
```

Assets CSV File Format

Assets represent individual nodes on the network, such as servers and routers. For more information, see ["The Network Model" on page 119](#).

You can define a set of assets in ESM by batch loading asset definitions from an Assets CSV file. Asset CSV files contain the columns listed in the table, below.

When an assets CSV file is selected for import, by default only the first fifteen rows of data are displayed in Select Column Headers for the Asset Data panel. However, when the data is

imported into the ArcSight Manager, all the rows are imported. For more information, see ["Increasing the Number of Displayed Rows" on page 143](#).

For the wizard to determine how to process the imported data, the type of each column must be specified. For more information, see ["Specifying CSV Column Types" on page 133](#).

When the Next button is clicked in the Summary of Data to Import panel, the asset data is imported into the ArcSight Manager. The new assets are created in the /All Assets/Site Assets group. For example, if an asset called DMZCorpEmailServer was specified in the imported assets CSV file, a new asset is created at the following URI: /All Assets/Site Assets/DMZCorpEmailServer. When imported, the new assets are auto-zoned. For more information, see ["Auto-Zoning Imported Assets" on page 160](#).

Assets CSV File Columns

Column Type	Description	Required Column?	Repeatable Column?	Example Value
Name	A descriptive name for the asset. This name must be unique. It is recommended to specify a name. However, if a name is not specified, a unique name is generated using the other fields.	No	No	DMZ Corp Email Server 1
Host Name	The host name of the network device represented by the asset.	No	No	dmz_corp_em11
IP Address	The IP address of the network device represented by the asset. NOTE: If no value is specified for this column (,) the asset is created with an IP address of 0.0.0.0.	Yes	No	192.0.2.0
MAC Address	The MAC address of the network device represented by the asset. The MAC address is made up of six groups of two hexadecimal digits can be separated by colons (:) or hyphens (-).	No	No	00-00-5E-00-53-00

Assets CSV File Columns, continued

Column Type	Description	Required Column?	Repeatable Column?	Example Value
Static Addressing	Defines if the network device is statically addressed even though the IP address of the asset is in a dynamic zone: • true—asset uses static IP addressing • false—device uses dynamic addressing (DHCP) For more information, see "Static Addressing in a Dynamic Zone" below .	No Default is false	No	false
Category URI	The asset category to assign to network device. NOTE: The wizard does not create new categories. For the category to be assigned, it must already exist.	No	Yes This column can be repeated because a network device can be categorized into more than one asset category.	/All Asset Categories/Site Asset Categories/Business Impact Analysis/Network Domains/Email/
Ignore	The column contains data that is not used by the Network Model Wizard when creating assets. For example, this column could contain a description of the asset.	No	Yes	This asset defines the Corporate Email Server in the DMZ.

An Example of an Assets CSV File

Here is an example of the Assets CSV file:

```
Lab Test machine,lab-111,<IP-address>,<Mac-address>,true,/All Asset  
Categories/ArcSight System Administration/Consoles/,/All Asset  
Categories/ArcSight System Administration/Databases/
```

Static Addressing in a Dynamic Zone

Set the **Static Addressing** column to true if the network device is statically addressed even though the IP address of the asset is in a dynamic zone. For example, set this column to true, for the following conditions:

- A dynamic zone is defined with an IP range, for example: 192.0.2.0 - 192.0.2.12.
- A network device with an IP address such as 192.0.2.5 is statically addressed even though it is defined in the dynamic zone.

For more about static and dynamic zones, see ["Zones" on page 124](#).

Asset Ranges CSV File Format

Asset ranges represent sets of network nodes addressable as a contiguous block of IP addresses. Asset ranges are useful when you have a number of network nodes that would be impractical to track individually, or that may come and go from the network, such as laptops. An asset range can define a group of assets that are not addressed individually. Asset ranges should be a subset of the IP address ranges defined for zones.



Caution: Each asset range should specify a unique range of IP addresses. The IP addresses specified by asset ranges should not overlap. If you import an asset range that overlaps with an asset range already specified on the ArcSight Manager and the new asset range has a different name than the existing asset range, the following occurs:

- The new asset range is created.
- The existing asset range is invalid and displays with the broken asset range icon in the .

You can define a set of asset ranges in ESM by batch loading asset range definitions from an asset range CSV file. Asset range CSV files contain the columns listed in the table, below. When an assets CSV file is selected for import, by default only the first fifteen rows of data are displayed in Select Column Headers for the Asset Ranges Data panel. However, when the data is imported into the ArcSight Manager, all the rows are imported. For more information, see ["Increasing the Number of Displayed Rows" on the next page](#).

For the wizard to determine how to process the imported data, the type of each column must be specified. For more information, see ["Specifying CSV Column Types" on page 133](#).

When the Next button is clicked in the Summary of Data to Import panel, the asset range data is imported into the ArcSight Manager. The new asset ranges are created in the /A11 Assets/Site Assets group. For example, if an asset range called DMZCorpHR was specified in the imported asset range CSV file, a new asset range is created at the following URI: /A11 Assets/Site Assets/DMZCorpHR.

Asset Range CSV File Columns

Column Type	Description	Required Column?	Repeatable Column?	Example Value
Name	A descriptive name for the asset range. This name must be unique.	Yes	No	DMZ Corp HR
Start Address	The start of the range of IP addresses that defines the asset range.	Yes	No	192.0.2.11

Asset Range CSV File Columns, continued

Column Type	Description	Required Column?	Repeatable Column?	Example Value
End Address	The end of the range of IP addresses that defines the asset range.	Yes	No	192.0.2.20
Category URI	The asset category to assign to asset range. NOTE: The wizard does not create new categories. For the category to be assigned, it must already exist.	No	Yes This column can be repeated because an asset range can be categorized into more than one asset category.	/All Asset Categories/Site Asset Categories/Business Impact Analysis/Data Role/HR Data/
Ignore	The column contains data that is not used by the Network Model Wizard when creating asset ranges. For example, this column could contain a description of the asset range.	No	Yes	This asset range defines the all the corporate human resources assets.

An Example of an Asset Ranges CSV File

Here is an example of the Asset Ranges CSV file:

```
HRRangeA,<Starting-IP-address>,<Ending-IP-address>,/All Asset  
Categories/ArcSight System Administration/Databases/
```

```
IT Range X,<Starting-IP-address>,<Ending-IP-address>,/All Asset  
Categories/ArcSight System Administration/Databases/
```

Increasing the Number of Displayed Rows

When the data is imported into the ArcSight Manager, all the rows are imported. However, by default, only the first fifteen rows of data are displayed in Select Column Headers for the <Resource Type> Data panels.

To increase the number of displayed rows, add the property `usecase.networkmodeling.maxrowfortable` to the <ARCSIGHT_HOME>/config/console.properties file and set the value of the property to a number greater than fifteen.

See "Managing and Changing Properties File Settings" in the [ESM Administrator's Guide](#) for procedures to edit properties files.

Summary of Data to Import

In the Summary of Data to Import panel, a summary of the network modeling data ready to import into the ArcSight Manager is displayed.

1. Click **Next** to start the import process.

A temporary Archive Resource Bundle (ARB) file with the import data is created and the Install Packages dialog appears.

2. To install the data from the temporary ARB file, click **OK** in the Update Packages dialog.

The network modeling data is imported into the ArcSight Manager and the Data Imported pane displays. In addition, the Installing Packages and the Importing Packages dialogs appear.

3. Close the Installing Packages and Importing Packages dialogs.

Network Data Imported into ArcSight Manager

When network modeling data is imported from the network modeling data CSV files, new resources are created in the following groups on the ArcSight Manager:

- New **zones** are created in the /All Zones/Site Zones group. For example, if a zone called DMZPublic was specified in the imported zones CSV file, a new zone is created at the following URI: /All Zones/Site Zones/DMZ Public.

The new zones are assigned to the default network called Local.

- New **assets** are created in the /All Assets/Site Assets group. For example, if an asset called DMZCorpEmailServer was specified in the imported assets CSV file, a new asset is created at the following URI: /All Assets/Site Assets/DMZCorpEmailServer. When imported, the new assets are auto-zoned. For more information, see ["Auto-Zoning Imported Assets" on page 160](#).
- New **asset ranges** are created in the /All Assets/Site Assets group. For example, if an asset range called DMZCorpHR was specified in the imported asset range CSV file, a new asset range is created at the following URI: /All Assets/Site Assets/DMZCorpHR.

In the Data Imported dialog, click **Finish** to close the wizard.

Working with Assets, Locations, Zones, Networks, Vulnerabilities, and Categories

The Assets resource provides tools for managing assets and asset ranges, and tools for managing the other network and asset modeling features associated with assets:

- Assets (["Managing Assets" below](#))
- Zones (["Managing Zones" on page 168](#))
- Networks (["Managing Networks" on page 170](#))
- Asset Categories (["Managing Asset Categories" on page 171](#))
- Vulnerabilities (["Managing Vulnerabilities" on page 162](#))
- Locations (["Managing Locations" on page 172](#))

Networks and *Zones* describe characteristics of how the asset is represented in the network itself. *Locations*, *Vulnerabilities*, and *Asset Categories* describe attributes of the assets that can be used for prioritization and correlation.

You can organize any of these distinctions into groups upon which you can set up user access controls.

You can also create a channel based on any of these distinctions to get additional monitoring views into the events happening on your network.

Managing Assets

This topic explains how to:

- Create, edit, move, and delete assets, and how to select them in the Common Conditions Editor. For an overview of what assets are, the resources that comprise them, how they fit into the network model, and the ways to populate the network model, see ["The Network Model" on page 119](#).
- Show assets in an active channel

Where: Navigator's Resources tab > Assets

To create or edit an asset:



Tip: In addition to creating assets manually using the (as described in this topic), you can create assets using the [Network Model wizard](#) or [dynamically from scanner data](#).

1. Select an asset group. Expand it if you are editing an asset in that group.
2. If you are creating an asset, right-click the group and choose **New Asset**.
If you are editing an asset, right-click the asset and choose **Edit Asset**.
3. Select the **Attributes** tab and enter or change values in the fields described below.

Asset Attributes	Description
Name	The asset's friendly name. This field can default to the asset's host name or IP address.
IP Address	The asset's IP address, in dotted-decimal notation.
MAC Address	The unique hardware ID for the network device.
Host Name	The asset's DNS name.
Location	The location resource for the asset. See "Locations" on page 126. • Select the location using the drop-down arrow. The arrow displays the Locations resource tree. • Alternatively, if the resource tree has too many subgroups to traverse and you know the location by name, see "Using the Advanced Selector While Editing Resources" on page 88.
Zone	The zone resource for the asset. See "Zones" on page 1052. • Use the drop-down arrow to select the zone. The arrow displays the Zones resource tree. • Alternatively, if the resource tree has too many subgroups to traverse and you know the zone by name, use the Advanced Selector button. See "Using the Advanced Selector While Editing Resources" on page 88.

4. Use the other tabs in the Asset Editor as necessary to add resources:

Asset View	Contents
Categories	Use the Add button on this tab to select network categories with which to associate the asset.
Alternate Interfaces	Use the Add button on this tab to select a second asset ID if this asset has an additional ID on another network. Alternate interfaces usually apply only to network boundary devices, such as bridges, that have two MAC addresses.
Vulnerabilities	Use the Add button on this tab to select certain vulnerabilities with which to associate the asset.
Notes	Use the text box and Save button on this tab to write and file additional information concerning the asset. For more information, see "Using Notes" on page 79

5. Click **OK**.

To move or copy an asset:

1. Drag and drop the asset into another group.
2. Choose one:

Move to move the asset

Copy to make a separate copy of the asset

Link to create a copy of the asset that is linked to the original asset. Changes to either the original or linked asset changes both instances of the asset. Deletion of either the original or linked asset deletes both instances of the asset.

To delete an asset:



Caution: Take care when deleting assets. Asset groups required for correct operation are locked; however, depending on your permissions, it is possible to delete the individual assets in those groups, such as the assets automatically created to track ArcSight components.

Do not delete ArcSight System Administration assets without consulting an ArcSight administrator.

1. Right-click an asset and choose **Delete Asset**.
2. Click **Yes** to confirm.

To show assets in a channel:

1. Right-click an asset or group of assets and choose **Show Assets**.

The assets are displayed in an active channel grid view.

2. If applicable, you can also show not only assets in the selected group but also all children in the group. To do so, right-click an asset group, and choose **Show Assets Recursively**.

To find an asset:

If you want to save time locating one asset in a potentially large set on the Navigator, use resource search. See ["Finding Resources" on page 704](#) for instructions.

Asset Auto-Creation

ESM automatically creates assets for components and, if applicable, for assets arriving from scan reports sent by vulnerability scanners via scanner SmartConnectors.

As a configuration option, you can also configure the creation of assets for devices reporting through SmartConnectors.

This section describes in detail how assets are automatically created from vulnerability scan reports and, if so configured, for devices reporting through SmartConnectors.

Topics include:

- ["Creating Assets from a Vulnerability Scan Report" below](#), including subtopics
- ["Creating Assets for SmartConnectors" on the next page](#), including subtopics
- ["Creating Assets for Network Devices" on page 151](#)
- ["Asset Names" on page 156](#)
- ["Changing the Default Naming Scheme" on page 157](#)

Creating Assets from a Vulnerability Scan Report

Assets are created from vulnerability scan reports differently for dynamic and static zones (for more about dynamic and static zones, see ["Zones" on page 124](#)).



Tip: Scanner reports list only information received through the scanner, whereas Asset Editors include the full list of both scanner data and vulnerability mappings stored in the system. So, the Editors might show more or different information than that shown in scanner reports.

Creating Assets from a Vulnerability Scan Report for Static Zones

For assets in static zones, you need at least an IP address or a host name to create an asset. Asset identifiers are looked up in the following order in a static zone:

MAC address > IP address > host name

The table below describes the action taken based on available IP and host name information.

Vulnerability Assets in Static Zones

Example	Action taken if no conflicts	Action taken if previous asset has similar information
IP=203.0.113.0 hostname=myhost	Asset created	Previous asset updated
IP=203.0.113.6 hostname=null	Asset created	Previous asset updated
IP=null hostname=myhost	Asset created	Previous asset updated
IP=null hostname=null	Asset not created. Either IP or host name is required.	No action taken. Either IP or host name is required.

Creating Assets from a Vulnerability Scan Report for Dynamic Zones

For assets in dynamic zones, assets are identified by IP address and host name and/or MAC address. Asset identifiers are looked up in the following order in a dynamic zone is:

MAC address > host name > IP address

By default, assets are not created in a dynamic zone if there is no host name present. The property set by default is:

```
scanner-event.dynamiczone.asset.nonidentifiable.create=false
```

Also by default, previous assets with similar information are discarded. This ensures that the network model is kept up to date with devices that are actively reporting events. The default property is set like this:

```
scanner-event.dynamiczone.asset.ipconflict.preserve=false
```

You can configure the system to create an asset if the asset has either an IP address or a host name, or to preserve previous assets with similar information by customizing settings in `server.properties`.

If you want to change the default behavior, refer to the *ESM Administrator's Guide*. Under the "Configuration Tasks" section, see the topic, "Asset Auto-Creation from Scanners in Dynamic Zones."

Creating Assets for SmartConnectors

ESM auto-creates assets for the ArcSight SmartConnectors connected to it. Creating assets for SmartConnectors is affected by the following conditions:

- ESM does not create an asset if there is no Network defined for the SmartConnector. This could happen if a SmartConnector is added incorrectly, or if an unforeseen condition occurs, such as a database corruption. If you do not specify a Network for the Connector during setup, ESM uses the default RFC1918 system zones.
- ESM does not create an asset unless the event is a base event: that is, an event generated by the device whose events the SmartConnector represents. For example, ESM creates an asset for the SmartConnector if the event is a firewall event, but it does not create an asset for the SmartConnector if the event is an ArcSight internal event, such as heartbeat events with the Manager.
- ESM does not create an asset for the SmartConnector if the Connector Asset Auto Creation Controller filter (at /All Filters/ArcSight System/Asset Auto Creation/Connector Asset Auto Creation Controller) is specially configured to exclude traffic from assets in this zone.

If the Connector Asset Auto Creation Controller filter is configured to exclude Connectors in a certain zone, such as a zone designated for VPN traffic that comes and goes from the network, then ESM does not create an asset for that Connector every time VPN traffic comes in from that Connector. This ensures that assets are not created unnecessarily.

For more about the Asset Auto Creation Auto Controller filter and how to configure it, see the [ArcSight Administration and ArcSight System Standard Content Guide](#). For information about an optional ArcSight Foundation, refer to the Standard Content Guide for that Foundation.

Creating Assets for SmartConnectors in Static Zones

By default, the SmartConnector auto-create function is enabled. For static zones, ESM needs both the IP address and host name to positively identify the SmartConnector asset. The table below shows the action taken when base events come in from devices involving that SmartConnector.

SmartConnector Assets in Static Zones

Example	Action taken if no previous SmartConnector	If a SINGLE enabled SmartConnector with the same IP address exists in the same zone	If multiple previous SmartConnectors with the same name exist in the zone
ip=198.51.100.0 hostname=myhost	Asset created	Existing asset is relocated to the zone assigned to the Connector. If there's an asset with the same name in the group, the new asset is renamed by adding ' _#', for example, asset_1.	No asset created.
ip=198.51.100.255 hostname=null	Asset not created. Both the IP address and host name are required.	Existing asset is relocated to the zone assigned to the Connector. If there's an asset with the same name in the group, it gets renamed.	No asset created.
ip=null hostname=myhost	Asset not created. Both the IP address and host name are required.	Asset not created. Both the IP address and host name are required.	No asset created.
ip=null hostname=null	Asset not created. Both the IP address and host name are required.	Asset not created. Both the IP address and host name are required.	No asset created.

Creating Assets for SmartConnectors in Dynamic Zones

By default, the SmartConnector auto-create function is enabled. For dynamic zones, it first looks for a MAC address, then a host name to positively identify the SmartConnector asset. The table below shows the action taken when base events come in involving that SmartConnector. (If multiple previous assets with the same name exist in the zone, no asset is created.)

SmartConnector Assets in Dynamic Zones		
Example	Action taken if no previous SmartConnector	If a SINGLE enabled asset with the same IP address exists in the same zone
ip=198.51.100.0 hostname=myhost mac=00005E0053FF	Asset created.	Existing asset is relocated to the zone assigned to the Connector. If there's an asset with the same name in the group, the new asset is renamed by adding '_#', for example, asset_1.
ip=198.51.100.20 hostname=myhost mac=null	Asset created.	Existing asset is relocated to the zone assigned to the Connector. If there's an asset with the same name in the group, it gets renamed.
ip=198.51.100.61 hostname=null mac=00005E0053AB	Asset not created. Both the IP address and host name are required.	Asset not created. Both the IP address and host name are required.
ip=198.51.100.96 hostname=null mac=null	Asset not created. Both the IP address and host name are required.	Asset not created. Both the IP address and host name are required.
ip=null hostname=myhost mac=null	Asset not created. Both the IP address and host name are required.	Asset not created. Both the IP address and host name are required.
ip=null hostname=null mac=00005E0053DB	Asset not created. Both the IP address and host name are required.	Asset not created. Both the IP address and host name are required.

Creating Assets for Network Devices

By default, the Manager also auto-creates assets for the network devices that originate the events. This feature can be configured during Manager setup using the Manager Setup Wizard.

Creating assets for devices is affected by the following conditions:

- The Manager does not create an asset if there is no Network defined for the SmartConnector. This could happen if a SmartConnector is added incorrectly, or if an unforeseen condition occurs, such as a database corruption. If you do not specify a Network for the Connector during setup, ESM uses the default RFC1918 system zones.
- The Manager does not create an asset unless the event is a base event: that is, an event generated by the device whose events the SmartConnector represents. For example, the Manager creates an asset for the device if the event is a firewall event, but it cannot create an asset for the device if the event is an ArcSight internal event, such as heartbeat events with the Manager.
- The Manager does not create an asset for the device if the Connector Asset Auto Creation Controller filter (at /All Filters/ArcSight System/Asset Auto Creation/Device Asset Auto Creation Controller) is specially configured to exclude traffic from assets in this zone.

If the Connector Asset Auto Creation Controller filter is configured to exclude events from Connectors in a certain zone, such as a zone designated for VPN traffic that comes and goes from the network, then the Manager does not create an asset for the device the Connector represents every time VPN traffic comes in from that Connector. This ensures that The Manager does not create unnecessary assets.

For more about the Connector Asset Auto Creation Auto Controller filter and how to configure it, see the [ArcSight Administration and ArcSight System Standard Content Guide](#). If you have an optional ArcSight Foundation, refer to the same guide for information on that Foundation.

Creating Assets for Network Devices in Static Zones

If you configured Manager setup to auto-creates assets for the network devices generating assets, it takes the following actions based on IP address and host name in static zones.

Network Device Assets in Static Zones

Example	Action taken if no previous device	Asset with same information exists on any zone related to SmartConnector
IP=192.0.2.0 hostname=myhost	Asset created.	Move asset to a new group. If there is already an asset with the same name, the previous one is renamed.
ip=192.0.2.255 hostname=null	Asset not created. Both IP address and host name are required.	Asset not created. Both IP address and host name are required.
ip=null hostname=myhost	Asset not created. Both IP address and host name are required.	Asset not created. Both IP address and host name are required.
ip=null hostname=null	Asset not created. Both IP address and host name are required.	Asset not created. Both IP address and host name are required.

Creating Assets for Network Devices in Dynamic Zones

If you configured Manager setup to auto-create assets for the network devices generating assets, it takes the following actions based on IP address, host name, and MAC address in dynamic zones.



Note: To auto-create assets for network devices in dynamic zones, both IP address and host names are required.

Network Device Assets in Dynamic Zones

Example	Action taken if no previous device	Asset with same information exists on any zone related to SmartConnector
ip=192.0.2.0 hostname=myhost mac=00005E005300	Asset created.	Move asset to a new group. If there is already an asset with the same name, the previous one is renamed.
ip=192.0.2.255 hostname=myhost mac=null	Asset created.	Move asset to a new group. If there is already an asset with the same name, the previous one is renamed.
ip=192.0.2.11 hostname=null mac=00005E0053AB	Asset not created. Host name is required.	Asset not created. Host name is required.
ip=192.0.2.30 hostname=null mac=null	Asset not created. Host name is required.	Asset not created. Host name is required.
ip=null hostname=myhost mac=null	Asset not created. IP address is required.	Asset not created. IP address is required.
ip=null hostname=null mac=00005E00FF	Asset not created. Both IP address and host name are required.	Asset not created. Both IP address and host name are required.

Asset Auto-Creation from Scanners in Dynamic Zones

The following properties relate to how the Manager creates assets from a vulnerability scan report for dynamic zones.

Create Asset with IP Address or Host Name

By default, the Manager does not create an asset in a dynamic zone if there is no host name present, as described in ["Creating Assets from a Vulnerability Scan Report for Dynamic Zones" on page 149](#). The property set by default is:

```
scanner-event.dynamiczone.asset.nonidentifiable.create=false
```

You can configure the Manager to create the asset as long as it has either an IP address or a host name. In `server.properties`, change `scanner-event.dynamiczone.asset.nonidentifiable.create` from **false** to **true**. The Manager discards conflicts between an IP address and host name (similar IP address, but different host name and/or MAC address).



Caution: Creating an asset if no host name is present can result in an inaccurate asset model.

Setting `scanner-event.dynamiczone.asset.nonidentifiable.create` to `true` means that assets are created if the asset has either an IP address or a host name.

This could lead to disabled assets or duplicated assets being created. Change this configuration only if you are using a dynamic zone to host ostensibly static assets, such as long-lived DHCP addresses.

When this property is set to `true`, the Manager takes the following actions.

Asset Creation with IP Address Only or Hostname Only

Example	Action taken if no conflicts	Action taken if previous asset with similar information
IP=198.51.100.0 hostname=myhost mac=00005E005300	Asset created.	Asset created, previous asset is deleted.
ip=198.51.100.60 hostname=myhost mac=null	Asset created.	Asset created, previous asset is deleted.
ip=198.51.100.100 hostname=null mac=00005E0053FF	Asset created.	Asset created, previous asset is deleted.
ip=198.51.100.255 hostname=null mac=null	Asset created.	Asset created, previous asset is deleted.

Asset Creation with IP Address Only or Hostname Only, continued

Example	Action taken if no conflicts	Action taken if previous asset with similar information
ip=null hostname=myhost mac=null	Asset not created.	
ip=null hostname=null mac=00005E0053AB	Asset not created.	
ip=null hostname=myhost mac=00005E005302	Asset not created.	

Preserve Previous Assets

This setting applies when the Manager creates assets from a vulnerability scan report for dynamic zones. By default, if a previous asset with similar information already exists in the asset model, the Manager creates a new asset and delete the old one, as described in ["Creating Assets from a Vulnerability Scan Report for Dynamic Zones" on page 149](#).

If you want to preserve the previous asset rather than delete it when a scan finds a new asset with similar information, you can configure the Manager to rename the previous asset. In `server.properties`, change `scanner-event.dynamiczone.asset.ipconflict.preserve` from **false** to **true**.



Caution: Preserving previous assets results in a larger asset model.

Setting `event.dynamiczone.asset.ipconflict.preserve` to **true** means that assets are continually added to the asset model and not removed. Use this option only if you know you must preserve all assets added to the asset model.

When you configure the Manager with `scanner-event.dynamiczone.asset.nonidentifiable.create=false` and `scanner-event.dynamiczone.asset.ipconflict.preserve=true`, it takes the following actions:

ESM Actions for Preserved Previous Assets

Example	Action taken if previous asset with similar information and preserve = true
IP=203.0.113.0 hostname=myhost mac=00005E005300	Asset created, previous asset is renamed.
ip=203.0.113.2 hostname=myhost mac=null	Asset created, previous asset is renamed.
ip=203.0.113.84 hostname=null mac=00005E0053FF	Asset created, previous asset is renamed.
ip=203.0.113.10 hostname=null mac=null	No action taken. Either host name or MAC address is required.
ip=null hostname=myhost mac=null	Asset not created.
ip=null hostname=null mac=00005E0053AB	Asset not created.
ip=null hostname='myhost' mac=00005E005302	Asset not created.

Asset Names

The Manager names the auto-created assets using the following templates. The creation rules work differently depending on how the events arrive: by Connector or by scanner; and whether they belong to a dynamic or static zone.

Asset Names from Scanner Events

By default, assets that come from scanners use the naming scheme outlined below, depending on whether the assets came from a static or dynamic zone. This scheme controls how asset names appear in channels and labels in the user interface.

Asset Names from Scanner Events

	Static Zone	Dynamic Zone
Property	scanner-event.auto-create.asset.name.template	scanner-event.auto-create.dynamiczone.asset.name.template
Value	\$destinationAddress - \$!destinationHostName	\$destinationHostName
Example	198.51.100.0 - myhost	myhost

You can reconfigure this default naming scheme the `server.properties` file, for example, if you want to show the host name first, or use an underscore to separate the elements. For details, see "Advanced Configuration for Asset Auto Creation" and "Changing the Default Naming Scheme" in the [ESM Administrator's Guide](#).

SmartConnector and Device Asset Names

SmartConnector and device assets are given the host name of the system that hosts them:

```
name = hostname
```

Changing the Default Naming Scheme

By default, the Manager names assets that come from scanners using the naming scheme outlined in ["Asset Names" on the previous page](#).

Asset Default Names		
	Static Zone	Dynamic Zone
Property:	scanner-event.auto-create.asset.name.template	scanner-event.auto-create.dynamiczone.asset.name.template
Value:	\$destinationAddress - \$!destinationHostName	\$destinationHostName
Example:	192.0.2.1 - myhost	myhost

You can reconfigure this default naming scheme, for example, if you want to show the host name first, or use an underscore to separate the elements.

For example, you want the asset name for an asset in a static zone to appear this way in the user interface:

```
myhost_192.0.2.1
```

In this case, change the default

```
$destinationAddress - $!destinationHostName
```

to

```
$!destinationHostName_$destinationAddress
```

Selecting Assets in the Common Conditions Editor


Purpose: After assets are added to your network model, you can select them when you are defining conditions that help you analyze the assets' role in the event being investigated.

Where used: The Asset Selector appears where you are defining conditions for these resources

- Rule's Conditions tab
See ["Adding Asset Conditions" on page 516](#) for an example.
- Query's Conditions tab
See ["Query Conditions" on page 333](#) for an example. The query is used in query viewers, trends, and reports.
- Filter's Filters tab
Same example as for rules.
- Active Channel's Filter tab
The channel is using a filter resource containing the asset condition.

Data Monitors use conditions by referencing a filter resource.

On the Common Conditions Editor panel, add an Asset condition by clicking the **Assets** on the icon menu or by right-clicking a node and selecting **New "Assets" Condition**. Below is an example of how a new Assets condition is defined on the filter resource's Filters tab:



Auto-Zoning an Asset

Purpose: To assign an asset or a group of assets into a network zone.

Where: Navigator > Assets > an asset group

- You can auto zone up to 1,000 assets at a time using the Navigator.
 - You cannot use Auto Zone to move locked assets.
1. Right-click an asset or group of assets and choose **Auto Zone**.
 2. In the Network Selector dialog, browse for the network that containing the zone with an IP address range that includes the asset.
 3. Select the network and click **OK**.

If a matching zone with an address range that includes the selected asset can be found in the network, the zone is assigned to the asset.

For example, a zone called DMZCorporate is defined in the Local network on the ArcSight Manager with a starting address of 192.0.2.0 and an ending address of 192.0.2.22. If an asset called DMZCorpDatabase with an IP address of 192.0.2.11 is selected for auto zoning in the Local network, the DMZCorporate zone is assigned to DMZCorpDatabase asset because the IP address of the DMZCorpDatabase asset is in the range of addresses specified in the DMZCorporate zone.

If no matching zone is found in the network, no zone is assigned.

Auto zoning can automatically occur when assets are imported using the Network Model Wizard. For more information, see ["Auto-Zoning Imported Assets" below](#).

Auto-Zoning Imported Assets

When new assets are imported into the ArcSight Manager using the Network Model Wizard, an attempt is made to assign the assets to the appropriate zone from the default network called Local. This process is called auto-zoning.

When the asset is imported, if a zone is found with an address range that includes the imported asset and that zone is located in the Local network, the matching zone is assigned to the asset. For the asset to find the matching zone, the matching zone must either:

- Already exist on the ArcSight Manager prior to the import.
- Be imported with the asset as part of the same import process—part of the same transaction. Zones are created before assets in the import process.

If no matching zone is found in the network, no zone is assigned.

The following example illustrates the auto-zone process. A zone called DMZCorporate is defined in the Local network on the Manager with a starting address of 192.0.2.0 and an ending address of 192.0.2.22. If an asset called DMZCorpDatabase with an IP address of 192.0.2.11 is imported by the wizard, the DMZCorporate zone is assigned to DMZCorpDatabase asset because the IP address of the DMZCorpDatabase asset is within the range of addresses specified in the DMZCorporate zone, and the DMZCorporate zone is located in the Local network.



Note: Only one asset with a given host name is allowed in a given zone on a network. When two assets with the same host name are imported, and if the Manager assigns them to the same zone in the same network, both assets are imported but one of the assets is disabled and displays with the broken-asset icon in the .

Managing Asset Groups

Asset groups are created to store similar groups or assets in a single location. Groups can be created within groups to meet enterprise needs. When a group is created within a group, the new group inherits the existing group's permissions. If a group is deleted, the assets within that group are also deleted. ArcSight provides these groups:

- **Shared:** this group lists assets to which the user has permission.
- **Unassigned:** this group lists assets not assigned to a group.

If you have Administrator access you also see another group named "All Assets" that contains all asset groups and assets.



Caution: Do not exceed 10,000 assets for each asset group. This is to ensure that automatic aging of assets works as expected. Asset aging is described in ["Asset Aging and Model Confidence" on page 123](#). For instructions on how to configure asset aging, refer to the "Asset Aging" topic in the [ESM Administrator's Guide](#).

To create an asset group:

1. In the Navigator panel's drop-down menu, choose **Assets**.
2. In the Assets resource tree, right-click a group and choose **New Group**. A "name" text field appears under the group you selected.
3. In the name text field, type in a name.
4. Press **Enter**.

To rename an asset group:

1. In the Assets resource tree, right-click a group and choose **Rename**.
2. In the "name" text field, rename the group.

To edit an asset group:

1. In the Assets resource tree, right-click a group and choose **Edit Group**.
2. In the **Group Editor**, edit the Name and Description text fields.

To move or copy an asset group:

1. In the Assets resource tree, navigate to a group and drag and drop it into another group.
2. Choose **Move** to move the group, **Copy** to make a separate copy of the group, or **Link** to create a copy of the group that is linked to the original group.

If you choose **Copy**, you create a separate copy of the group that is not affected when the original group is edited. If you choose **Link**, you create a copy of the group that is linked to the original group. Therefore, if you edit a linked group, whether the original or the copy, all links are edited as well. When deleting linked groups, you can either delete the selected group or all linked groups.

To delete an asset group:

1. In the Assets resource tree, right-click a group and choose **Delete Group**.
2. In the dialog box, click **Yes**.

Managing Vulnerabilities

This topic describes how to perform the authoring and management tasks for vulnerabilities such as creating, editing, moving, and retrieving vulnerable assets.

See also ["Modeling the Network" on page 119](#).

Note also that you can create a vulnerability channel. For more information on active channels, see ["Monitoring Active Channels" on page 224](#).

Where: Navigator > Resources > Assets > Vulnerabilities tab

To create a vulnerability:

1. In the Navigator panel's drop-down menu, choose **Assets**, then click the **Vulnerabilities** tab.
2. Right-click a group and choose **New Vulnerability**.
3. On the Vulnerabilities Attributes tab, type in the following text fields:

Vulnerability Attribute	Description
Name	The vulnerability's name (required). It can be generated by the ArcSight Manager in response to vulnerability scanners. If so, this field is identical to the External ID field except that the pipe () is replaced with a dash (-). For example, CVE CVE-1999-200 is represented as CVE - CVE-1999-200.
Knowledge Base Article	Optional: A link to a knowledge base article that further describes the vulnerability.
External ID	An ID of the format <standards body> <id>, such as CVE CVE-1999-200.
Owners	ArcSight users (analysts) who are interested in the vulnerability.
Notification Groups	ArcSight users (analysts) who are notified of events involving the vulnerability.

4. On the Vulnerable Assets tab, click the **Add New** button, if you've defined assets that include this vulnerability.



Note: Refer to ["Working with Vulnerable Assets"](#) on page 165 for details on using the Vulnerable Assets tab.

To edit a vulnerability:

1. Right-click a vulnerability and choose **Edit Vulnerability**.
2. On the Attributes tab, type in the text fields as described above.
3. On the Vulnerable Assets tab, click the **Add New** button, if you've defined assets that include this vulnerability.

To move or copy a vulnerability:

1. Drag and drop a vulnerability into another group.
2. Choose one:
 - **Move** to move the vulnerability,
 - **Copy** to make a separate copy of the vulnerability, or
 - **Link** to create a copy of the vulnerability that is linked to the original vulnerability.

If you choose **Copy**, you create a separate copy of the vulnerability that is not affected when the original vulnerability is edited. If you choose **Link**, you create a copy of the vulnerability that is linked to the original vulnerability. Therefore, if you edit a linked vulnerability, whether it be the original or the copy, all links are edited as well. When deleting linked vulnerabilities, you can either delete the selected vulnerability or all linked vulnerability copies.

To delete a vulnerability:

1. Right-click a vulnerability and choose **Delete Vulnerability**.
2. In the dialog box, click **Yes**.

To add a vulnerability to an asset:

1. Open a vulnerability active channel
2. Right-click a vulnerability and choose **Add To Asset**.
3. In the Asset Editor, click **OK**.

Selecting Vulnerabilities in the Common Conditions Editor

Purpose: After assets are added to your network model, you can select them while you are defining conditions to track associated vulnerabilities.

Where used: The Vulnerabilities selector appears when you are adding a Has Vulnerability condition to an asset.

- Rule's Conditions tab

See ["Adding Vulnerability Conditions" on page 516](#)

- Query's Conditions tab

See ["Query Conditions" on page 333](#) for an example. The query is used in query viewers, trends, and reports.

- Filter's Filters tab


Same example as for rules.

- Active Channel's Filter tab

The channel is using a filter resource containing the asset condition.

Data Monitors use conditions by referencing a filter resource.

On the Common Conditions Editor (CCE) panel, add a Has Vulnerability condition on an asset node and selecting **New "Has Vulnerability" Condition**:



Working with Vulnerable Assets

This topic describes tasks associated with the vulnerability resource's Vulnerable Assets tab.

Where: Navigator > Resources > Assets > Vulnerabilities tab

To retrieve vulnerable assets:

1. Right-click a vulnerability and choose **Edit Vulnerability**.
2. Select the **Vulnerable Assets** tab.

If you used a vulnerability scanner, all vulnerable asset discovered by the scanner are listed on this tab.

3. To refresh the vulnerabilities list, click the **Refresh** button ().

To add an asset to a vulnerability list:

1. Right-click a vulnerability and choose **Edit Vulnerability**.
2. In the Vulnerability Editor, select the **Vulnerable Assets** tab.

3. Click the **Add** button ().
4. Select an asset in the Assets Selector and click **OK**.

To delete an asset from a vulnerability list:

1. Right-click a vulnerability and choose **Edit Vulnerability**.
2. In the Vulnerability Editor, select the **Vulnerable Assets** tab.
3. Select an asset and click the **Delete** button ().
4. In the dialog box, click **Yes**.

Managing Vulnerability Groups

This topic describes the tasks involved in managing vulnerability groups.

Where: Navigator > Resources > Assets > Vulnerabilities tab

To create a vulnerability group:

1. Right-click a vulnerability group and choose **New Group**.
A "name" text field appears under the group you selected.
2. In the "name" text field, type in a name.
3. Press **Enter**.

To rename a vulnerability group:

1. Right-click a vulnerability group and choose **Rename**.
2. In the "name" text field, rename the group.
3. Press **Enter**.

To edit a vulnerability group:

1. Right-click a vulnerability group and choose **Edit Group**.
2. In the Group Editor, edit the **Name** and **Description** text fields.
3. Click **OK**.

To move or copy a vulnerability group:

1. Navigate to a vulnerability group and drag and drop it into another group.
2. Choose:

- **Move** to move the group,
- **Copy** to make a separate copy of the group, or
- **Link** to create a copy of the group that is linked to the original group.

If you choose **Copy**, you create a separate copy of the group that is not affected when the original group is edited. If you choose **Link**, you create a copy of the group that is linked to the original group. Therefore, if you edit a linked group, whether the original or the copy, all links are edited as well. When deleting linked groups, you can either delete the selected group or all linked groups.

To delete a vulnerability group:

1. Right-click a vulnerability group and choose **Delete Group**.
2. In the dialog box, click **Yes**.

Showing Affected Assets

The **Analyze in Channel** options on the grid view include the ability to explore events that potentially exploit asset vulnerabilities. You can also view an event's targeted assets.

To show exploited vulnerabilities:

1. Select an event in a grid view.
2. Right-click the event and choose **Analyze in Channel > Show Exploited Vulnerabilities**. Available information appears in the Vulnerabilities tab of the relevant Asset Editor.

To show an event's targeted asset:

1. Select an event in a grid view.
2. Right-click the event and choose **Analyze in Channel > Show Targeted Asset**. Available information appears in the Asset Editor.

Assets are part of your network model. Refer to ["Modeling the Network" on page 119](#) for more information.

Reporting on Output from Vulnerability Scanners

You can review the output of asset-vulnerability scanners in active channels and in the Vulnerabilities tab of the Asset Editor.

1. Choose the Assets resource tree in the Navigator panel.
2. In the Assets tab of the Assets tree, right-click an individual asset and choose **Scanner Reports**. If scanner asset-vulnerability reports are available for the selected asset, they appear in a Viewer panel grid view as an active channel.
3. You can use the standard controls described in Using Grids and Active Channels to review the reports collectively.
4. Also in the channel view, you can double-click vulnerability scanner events to open them in the Asset Editor, where the Vulnerabilities tab lists the vulnerability details.

For information on creating and editing assets, see ["Modeling the Network" on page 119](#).

You can create an active channel for selected scanner reports. For information on using active channels, see ["Monitoring Active Channels" on page 224](#).

Reporting on Asset Vulnerabilities

You can create reports to show which assets are vulnerable to particular vulnerabilities or threats. ArcSight also provides Asset Reports that can be run from the Reports resource tree in the Navigator panel. For more information, see ["Using Report Templates" on page 426](#).

1. In the Navigator panel's drop-down menu, choose **Assets**, then click the **Vulnerabilities** tab.
2. In the Vulnerabilities tree, right-click a vulnerability and choose **Vulnerable Assets Report**.
3. In the Report Parameters dialog box, accept the vulnerability listed in the **Vulnerability URL** text field or click the **Vulnerability** button to run the report on another vulnerability.
4. Choose a Report File Format from the drop-down menu and click **OK**.

Reports can be archived in PDF, HTML, Excel, Comma Separated Value (csv), or Rich Text Format (rtf). The default PDF format should be used when archiving reports. Compared to PDF reports, other reports may lose formatting information and look different from what you expect. In addition, Excel format is more memory-intensive than PDF.

Managing Zones

For an overview of zones and how they fit into the network model, see ["Zones" on page 124](#).

Shrinking or splitting zones

- The Zone Editor cannot be used to shrink a zone if there are assets that fall outside the range of the new zone. For example, if you have a zone with an address range of 192.0.2.1 to 192.0.2.27 and an asset in that zone with an IP address of 192.0.2.15, you cannot change the upper end of the zone range to 192.0.2.10 but you can change it to 192.0.2.20.
- For shrinking or splitting zones that might encounter such issues, we suggest using a package export and import operation. You can export the asset resources and then import them back in. Package import and install automatically assigns assets to appropriate zones similar to the *auto-zoning* used by the Network Model Wizard. See ["Managing Packages" on page 730](#), ["Populating the Network Model Using the Wizard" on page 132](#), and ["Auto-Zoning Imported Assets" on page 160](#).

Keeping separate zones for mixed-family IP addresses

If your network has both IPv4 and IPv6 addresses, have separate zones for IPv4 and IPv6 addresses. When an IPv4 address is compared to an IPv6 address or used in a mixed-family zone, the IPv4 address is handled as IPv4-mapped IPv6. If you create a mixed-family zone, some users may find results of this mapping as unexpected. OpenText therefore recommends creating zones that just contain one type of address.

Zone Attributes

Attribute	Description
Name	A descriptive name for the IP address range the network zone represents (required)
Start Address	Provide an IP address that identifies the start of the network scope.
End Address	Provide an IP address that identifies the end of the network scope.

Zone Attributes, continued

Attribute	Description
Dynamic Addressing	Click this option on or off to indicate whether this network uses dynamic addressing • Checkmark (toggle on) this option to indicate that the network you are describing uses dynamic addressing (Dynamic Host Configuration Protocol or DHCP server) • Leave this option unchecked (toggle off), if the network you are describing does not use dynamic addressing (but, rather, uses static IP addresses)
Location	The location resource for this zone. See "Locations" on page 126. • Select the location using the drop-down arrow. The arrow displays the Locations resource tree. • Alternatively, if the resource tree has too many subgroups to traverse and you know the location by name, see "Using the Advanced Selector While Editing Resources" on page 88.
Network	The network in which this zone resides. • Select the zone using the drop-down arrow. The arrow displays the Zones resource tree. • Alternatively, if the resource tree has too many subgroups to traverse and you know the zone by name, see "Using the Advanced Selector While Editing Resources" on page 88.

In addition to the above zone **Attributes**, the Zone Editor includes subtabs for adding **Assets** and **Categories** into the *zone* you are configuring.

Managing Networks

For an overview of networks and how they fit into the network model, see ["Networks" on page 126](#).

Network Attribute	Description
Name	A descriptive name for the network (required)
Customer	The Customer name is typically used if configuring assets for a customer on behalf of a managed security service provider (MSSP). • Select the customer using the drop-down arrow. The arrow displays the Customers resource tree. • Alternatively, if the resource tree has too many subgroups to traverse and you know the customer by name, use the Advanced Selector  button. See "Using the Advanced Selector While Editing Resources" on page 88.
Location	This is an optional field for a descriptive name of the geographical location of the network. • Select the location using the drop-down arrow. The arrow displays the Locations resource tree. • Alternatively, if the resource tree has too many subgroups to traverse and you know the location by name, use the Advanced Selector  button. See "Using the Advanced Selector While Editing Resources" on page 88.

In addition to network **Attributes** (described above), the Network Editor includes subtabs for adding **Connectors** and **Zones** into the selected *network* you are configuring.

Managing Asset Categories

To view the available asset categories:

1. On the Navigator resource tree, choose **Assets**.
2. Go to the **Categories** tab and expand a node, for example, **Site Asset Categories**.

The Categories tab provides options to organize assets into groups based on *categories*.

From the Navigator right-click menu on **Asset Categories**, you have several views and tools to help manage assets. From this menu, you can:

- Create channels to show asset categories and assets.
 - **Show Asset Categories** displays all top-level asset categories of an asset category group.
 - **Show Asset Categories Recursively** displays sublevels of asset categories, if any, of an asset category group
- Move assets into and out of category groups.
- Create new category groups.
- Configure access control lists (ACLs) to limit or allow user access to groups of assets (see ["Managing Permissions" on page 107](#)).

One asset can have multiple asset categories. You can also assign asset categories to groups of resources. This transfers the asset category onto all the members of the group and its subgroups.

To assign an asset category:

1. In the Navigator drop-down menu, go to **Assets**. Select the **Assets** tab. Go to ArcSight System Administration/Agents, to find the SmartConnectors installed for your environment.
2. Right-click the asset or asset group you wish to categorize and select **Edit Asset** (or **Edit Group**).
3. In the Inspect/Edit panel, click the **Categories** tab. Click the add icon () at the top of the screen.
4. In the Asset Categories Selector pop-up window, select the asset categories that apply to this asset and click **OK**. For example:

- a. The usage category that applies to the asset (for example, /All Asset Categories/Site Asset Categories/Business Impact Analysis/Business Role/Revenue Generation)
 - b. The criticality level that applies to the asset (for example, /All Asset Categories/System Asset Categories/Criticality/Very High)
5. Repeat steps 3 and 4 for every asset or group of assets you want to classify in one of the asset categories.

For an overview of asset categories and how they fit into the network model, see ["Asset Categories" on page 127](#).

Managing Locations

For an overview of locations and how they fit into the network model, see ["Locations" on page 126](#).

Where: Navigator > Assets > Locations subtab

To create or edit a location:

1. Right-click an unlocked Locations group and choose **New Location**.
2. Set the following attributes:

Location Attribute	Description
Name	A descriptive name for the geographical location (required)
Latitude	Latitude for the location. The format for this measurement is a preference setting for the (menu option Edit > Preferences , click Latitude and Longitude). For more information, see "Setting Latitude and Longitude Options" on page 56 .
Longitude	Longitude for the location The format for this measurement is a preference setting for the (menu option Edit > Preferences , click Latitude and Longitude). For more information, see "Setting Latitude and Longitude Options" on page 56 .
Address	Provide details for City , Region Code , Postal Code , and Country as required.

Managing Customers

Purpose: Customer tagging is a feature developed mainly to support Managed Security Services Provider (MSSP) environments, although it can also be used by private organizations

to denote cost centers, internal groups, or subdivisions. The Customer designation keeps event traffic from multiple cost centers and/or business units clearly identified and separate. For more general information about this feature, see *ESM 101*.

The Customers resource tree, when populated, maps out the various external or internal customer accounts your enterprise tracks for cost, security analysis, or administrative reasons. These accounts, if present, are usually set up as part of the ArcSight deployment process. If the Customers resource tree is abbreviated or empty, your organization is probably not using this feature.

When the Customers resource tree is populated, you primarily use its branches as references in analysis filters that exclude or include certain customers.

Apart from analysis, the activities necessary to maintain the Customers resource tree include creating new customer references, editing existing references, and occasionally deleting references.

Where: Navigator > Resources > Customers

To create customers:

When you create a customer, remember that the branch you add to the resource tree has to **match** the Customer URI attribute configured for that branch in the relevant SmartConnectors. In other words, you create customer-tracking resources only for those customers that have parallel URI values set in the SmartConnectors that monitor their devices.

1. Right-click a customer group and choose **New Customer**.
2. In the Customer Editor, enter values for the properties that identify the customer. Note that the **Name** value has to complete the correct Customer URI for this account as found in its related SmartConnectors.
3. On the **Networks** tab, add the network for this customer, as required.
4. Click **Apply** to update the customer and leave the editor open, or **OK** to complete editing and close the editor.

To edit customers:

1. Right-click a customer and choose **Edit Customer**.
2. Change the values, as appropriate.
3. Click **Apply** to update the customer and leave the editor open, or **OK** to complete editing and close the editor.

To delete customers:

1. Right-click a customer and choose **Delete Customer**.
2. Click **Yes** to confirm the deletion.

Chapter 6: Managing SmartConnectors

ArcSight SmartConnectors, if installed and configured to send events to the Manager, are available on the for further configuration. SmartConnectors that appear on the Console as resources are called registered connectors.

You can configure ArcSight SmartConnectors to optimize their performance and increase their functionality. For example, you can configure SmartConnectors to enable aggregation, batching, and time filter correction functionality. You can also send control commands from the ArcSight Console or the ArcSight Command Center to SmartConnectors to manage the flow of events.

Selecting and Setting SmartConnector Parameters

From the ArcSight Console, use the Connector Editor to control SmartConnectors that are registered in the ArcSight Manager.

Configuring the SmartConnector

A SmartConnector can have a default and a number of alternate configurations.

An alternate configuration is a set of runtime parameters that is used instead of the default configuration during a specified portion of every day. For example, you might want to specify different batching schemes (by severity or size) for different times of a day. You can define more than one alternate configuration per destination and apply them to the destination for different time ranges during the day. For example, you can define a configuration for 8 am to 5 p.m. time range and another configuration for the 5 pm to 8 am time range.

Where: Navigator > Resources > Connectors

To define default configurations:

1. Right-click the SmartConnector you want to manage and choose **Configure**.
This opens the **Inspect/Edit** panel for the **Connector Editor**. On the Connector tab, the Name field is automatically populated with the name assigned during SmartConnector installation.
2. Refer to ["Connector Tab Configuration Fields" on page 177](#) for descriptions of the attributes.

3. On the **Default** tab, change any additional Batching, Time Correction or other parameters as desired, using the configuration fields explanations provided in ["Default Content Tab Configuration Fields" on page 178](#).

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

The Description entry associated with the setting provides tool tip information. These parameters are not localized since they come directly from the connector and the connector may contain new resources (since it may be a newer version).

4. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
5. Click **Apply** to add your changes and to keep the Connector Editor open, or click **OK** to close the editor.

The framework for connector commands operates in a similar way. Configuration of the connector command menu is achieved by sending the list of commands that are supported on the connector at registration time.

To create alternate configurations:

1. Open the Inspect/Edit panel of the SmartConnector.
2. On the **Default** tab, click **Add Alternate**.

A new tab, Alternate #1, is added to the edit panel. The alternate tab provides fields for entering a time interval.

3. Under **Time Interval**, enter times for **From** and **To**.

The configuration fields for Alternate #1 are initially the same as the Default's entries. Make additional changes as required, then click **Apply**.

4. Repeat the process if you want additional alternates using different time intervals and different parameters. For example, create alternates if you want varying batching schemes based on severity or size on certain times of the day.

If the time ranges of the combined alternate configurations do not span 24 hours, the default parameters will be used to cover the time intervals not already defined in the alternates.

Connector Editor Tabs

The Connector Editor provides several options you can adjust. The options are organized in tabs and subtabs, described below.

Tabs on the Connector Editor	
Connector Tab Configuration Fields	Basic identification, ownership, and date/time parameters.
Networks	The ArcSight networks to which the connector is or can be assigned. Click the Add button on this tab to assign the connector to a network. Note: Network and zone assignment for connectors in a tiered ESM architecture In a tiered ESM architecture, you are sending events from one ESM source to the ESM destination. The ESM installations are similar and therefore the networks and zones are the same. You must therefore assign connectors (including Forwarding Connector) to the same network, and the network must be assigned to the zone. This ensures the preservation of Reference Pages information, if this resource group attribute is configured. See the following topics: • "Working with Assets, Locations, Zones, Networks, Vulnerabilities, and Categories" on page 145 for information on how to configure zones and networks in your asset model • "Working with Resource Groups" on page 711 for information about the Reference Pages attribute
Default Content Tab Configuration Fields	Includes options for report batching, aggregation, and time corrections.
Default: Filters	A filter condition editor for constraining what the connector reports. (See "Managing SmartConnector Filter Conditions" on page 192 and "Common Conditions Editor (CCE)" on page 887 for details on how to define filters for connectors.)
Alternate: Content	A set of options identical to those under Default, which you can use to create alternate configurations. Appears only after you add an Alternate tab.
Alternate: Filters	A filter condition editor for constraining what the connector reports, in an alternate configuration.
Notes: Table	A text editor for, and tabular list of, configuration notes.
Notes: List	A text editor for, and text presentation of, configuration notes.

Connector Tab Configuration Fields

You do basic configuration through the Connector and Default: Content tabs. Many of these fields correspond to resource editor fields. See also ["Common Resource Attribute Fields" on page 725](#).

Fields on the Connector Tab

Field Name	Description
Name	The Name text field is automatically populated with the name assigned during SmartConnector installation.
ID	The identification string assigned during SmartConnector installation.
Status	The SmartConnector's current mode of operation.
Connector Location	A description of the (usually) physical location of the SmartConnector. This appears in all the events issued from the connector.
Device Location	A description of the (usually) physical location of the device the SmartConnector is monitoring. This appears in all the events issued from the connector.
Version	The connector's software version number.
Comment	Optional descriptive text
Model Import User	If this is a model import connector, select the associated user.

Default Content Tab Configuration Fields

After configuring SmartConnector to send events, you can configure their operation further through the settings listed in Modify Destination Settings section. The details for the selections are provided in the following tables.

The following table shows the configurable settings.

Configurable Settings

Name Field	Value Field
Batching	SmartConnectors can batch events to increase performance and optimize network bandwidth. When activated, SmartConnectors create blocks of events and send them when they either (1) reach a certain size or (2) the time window expires. You can also prioritize batches by severity, forcing the SmartConnector to send the highest-severity event batches first and the lowest-severity event batches later.
Enable Batching (per event)	Create batches of events of this specified size (50, 100, 200, 300, 400, 500, or 600 events). The default is 100. Caution: You could potentially lose data with batch sizes 500 and 600. Contact Customer Support before using 500 or 600 batch size.
Enable Batching (in seconds)	The SmartConnector sends the events if this time window expires (1, 5, 10, 15, 30, 60).
Batch By	This is Time Based if the SmartConnector should send batches as they arrive (the default) or Severity Based if the SmartConnector should send batches based on severity (batches of Highest Severity events sent first).

Configurable Settings, continued

Name Field	Value Field
Time Correction	The settings in this group provide several ways to fix problems with devices that do not report the time correctly.
Use Connector Time as Device Time	(No Yes) Override the time the device reports and instead use the time at which the connector received the event. This option assumes that the connector is more likely to report the correct time. The default is No .
Enable Device Time Correction (in seconds)	The SmartConnector can adjust the time reported by the device Detect Time, using this setting. This is useful when a remote device's clock isn't synchronized with the ArcSight Manager. This should be a temporary setting. The recommended way to synchronize clocks between Manager and devices is the NTP protocol.
Enable Connector Time Correction (in seconds)	The SmartConnector can also adjust the time reported by the Connector Time SmartConnector itself, using this setting. This is for informational purposes only and allows you to modify the local time on the SmartConnector. This should be a temporary setting. The recommended way to synchronize clocks between Manager and SmartConnectors is the NTP protocol.
Set Device Time Zone To	(Disabled <TimeZone>) (Default is Disabled) Ordinarily, it is presumed that the original device is reporting its time zone along with its time. And if not, it is then presumed that the SmartConnector is doing so. If this is not true, or the device isn't reporting correctly, you can switch this option from Disabled to GMT or to a particular world time zone. That zone is then applied to the time reported. Refer to the <i>SmartConnector User's Guide</i> for the updated list of zones.
Device Time Auto-correction	These are the time span and frequency factors for doing device-time auto-correction.
Future Threshold	The connector auto-corrects if the detect time is greater than the connector time by Future Threshold seconds.
Past Threshold	The connector auto-corrects if the detect time is earlier than the connector time by Past Threshold seconds.
Device List	A comma-separated list of the devices to which the thresholds apply. The default, (ALL) means all devices.
Time Checking	
Future Threshold	The number of seconds by which to extend the connector's forward threshold for time checking. Default is 5 minutes (300 seconds).
Past Threshold	The number of seconds by which to extend the connector's rear threshold for time checking. Default is 1 hour (3600 seconds).
Frequency	The SmartConnector checks its future and past thresholds at intervals specified by this number of seconds. Default is 1 minute (60 seconds).
Cache	Changing these settings does not affect the events cached, it only affects new events sent to the cache.

Configurable Settings, continued

Name Field	Value Field
Cache Size	SmartConnectors use a compressed disk cache to hold large volumes of events when the ArcSight Manager is down or when the SmartConnector receives bursts of events. This parameter specifies the disk space to use. The default is 1 GB which, depending on the connector, can hold about 15 million events, but it also can go down to 5 MB . When this disk space is full, the SmartConnector drops the oldest events to free up disk cache space. (5 MB, 50 MB, 100 MB, 200 MB, 250 MB, 500 MB, 1 GB, 2.5 GB, 5 GB, 10 GB, 50 GB.)
Notification Threshold	The size of the cache's contents at which to trigger a notification. Default is 10,000.
Notification Frequency	How often to send notifications when the Notification Threshold is reached. (1 min, 5 min, 10 min, 30 min, 60 min.)
Network	
Heartbeat Frequency	This setting controls how often the connector sends a heartbeat message to the ArcSight Manager. The default is 5 seconds , but it can go from 5 seconds to 10 minutes . Note that the heartbeat is also used to communicate with the SmartConnector; therefore, if its frequency is set to 10 minutes , then it could take as much as 10 minutes to send any configuration information or commands back to the SmartConnector.
Enable Name Resolution	(No Source/Dest only Yes) The SmartConnector tries to resolve IP addresses to host names, and host names to IP addresses, if the event rate allows it and if required. This setting controls this functionality. The Source, Target and Device IP addresses and Hostnames may also be affected by this setting. The Source/Dest Only choice means that the device address and device host name fields are ignored for name resolution. (The default is Yes .)
IPv6 Name Resolution	• IPv4 Only for Legacy Events (default) • IPv6 (Prefer IPv4 for reverse resolution) for Legacy Events • IPv6 (Prefer IPv6 for reverse resolution) for Legacy Events
Name Resolution TTL (secs)	This is the amount of time (Time to Live) the name resolution is to be in effect. The name resolution entries are cached for this time (default is 3600).
Wait For Name Resolution	(Yes No) If set to Yes, the SmartConnector waits for name resolution to be completed. When Yes is selected, event processing might be slowed down significantly and even cause lost events. (default is No).
Name Resolution Host Name Options	• Set host name only (default) • Set host name only (lowercase) • Set host and domain names • Set host and domain names (lowercase) For reverse resolution (IP Address to Host name), only the host name field is set. If host name only is not used, the host name is split up and put into both the DNS domain and the host name fields. This affects the source, destination, device and agent address.

Configurable Settings, continued

Name Field	Value Field
Name Resolution Domain from Email	(Yes No) If set to Yes , the host name and DNS domain fields are empty, and the corresponding user name field appears as an e-mail address, then the domain from the e-mail address is put in the DNS domain field. This only affects the source and destination fields (default is Yes).
Clear Host Names Same as IP Address	(Yes No) If set to Yes and the host name field is set to an IP Address that matches the corresponding IP Address field, then the host name field is cleared. This affects the source, destination, and device fields (default is Yes).
Set Host Names to IP Addresses When Unknown	(Yes No) If set to Yes , host names that remain unresolved are set to IP addresses (default is No).
Don't Resolve Host Names Matching	By default, host names are resolved to their IP addresses. You have the option to specify a regular expression for all or part of a host name <i>for which you do not want the system to attempt host name resolution to an IP address</i>. When this option is configured, the system cannot resolve host names matching this expression.
Don't Reverse-Resolve IP Ranges	By default, IP addresses are resolved to their domain names. You have the option to specify IP address ranges <i>for which you do not want the system to attempt reverse-resolution to domain names</i>. Click in the field to enter the IP address range. To enter a single IP address, enter the address under the From column and leave the To column blank, then click Apply. For an address range, enter the starting IP address under From and the ending address under To, then click Apply. This field allows you to enter a list of ranges. When this option is configured, the system cannot reverse-resolve IP addresses that fall within any of the specified ranges.
Remove Unresolvable Names/IPs from Cache	(Yes Yes (w/ negative cache) No) If set to No , unresolvable host names or IP addresses continue to be in the cache. If set to Yes , unresolvable host names or IP addresses are removed from the cache. If set to Yes (w/negative cache) , the connector remembers what names/IPs have been unresolvable so that time is not wasted trying to resolve them frequently.
Limit Bandwidth To	Default is Disabled . A list of bandwidth options you can use to constrain the connector's output over the network. (Disabled , 1 kbit/sec to 10 Mbits/sec.)
Transport Mode	You can configure the SmartConnector to cache to disk all the processed events it receives. This is equivalent to pausing the SmartConnector. However, you can use this setting to delay event-sending during particular time periods. For example, you could use this setting to cache events during the day and send them at night. You can also set the connector to cache all events, except for those marked with a very-high severity, during business hours, and send the rest at night. (Normal Cache Cache but send Very High severity events).

Configurable Settings, continued

Name Field	Value Field
Cache Mode	(Normal Drop if Dest Down) This option is meant to be used on a primary destination to control the caching behavior of the primary destination when it is down, and the connector starts sending events to the failover destination. In the Normal mode, events are cached and sent to the primary destination when it comes back up. In the Drop if Dest Down mode, the events are not cached and dropped and therefore not sent to the primary destination when it becomes available again (default is Normal).
Address-Based Zone Population Defaults Enabled	(Yes No) If Yes, the default zones built into the connector will be used to assign zones. These zones are only used if a network model has not been sent by ESM or ArcMC, or if that network model does not cover some addresses. If the Address-Based Zone Population setting (below) is specified, you may want to change this to No.
Address-Based Zone Population	If specified in setup or ArcMC, this is a comma-separated list that must contain a multiple of three items. The first of each three is the starting IP address of a zone, the second is the ending IP address of the zone, and the third is the URI of the zone to assign to addresses in that range. These zones are only used if a network model has not been sent by ESM or ArcMC, or if that network model does not cover some addresses. If Address-Based Zone Population Defaults Enabled is set to Yes, the zones specified here take precedence over those. For example for two zones this could be: 15.0.0.0,15.255.255.255,/All Zones/ArcSight System/Public Address Space Zones/Hewlett-Packard Company,17.0.0.0,17.255.255.255,/All Zones/ArcSight System/Public Address Space Zones/Apple Computer Inc.
Zone Population Mode	(Normal Rezone (override) No Zoning (clear)) Setting to Normal means zones are computed and assigned, if not already set. Rezone (override) re-computes and re-assigns already populated zones. No Zoning (clear) clears the zones, if already populated. (default is Normal).
Customer URI	Applies the given customer URI to events emanating from the connector. Provided the customer resource exists, all customer fields are populated on the ArcSight Manager. If this particular connector is reporting data that might apply to more than one customer, you can use Velocity templates in this field to conditionally identify those customers.

Configurable Settings, continued

Name Field	Value Field
Field Based Aggregation	This feature is an extension of basic connector aggregation. Basic aggregation aggregates two events if, and only if, the fields of the two events are the same per the following fields: • Agent ID • Name • Device event category • Agent severity • Destination address • Destination user ID • Destination port • Request URL • Source address • Source user ID • Source port • Destination process name • Transport protocol • Application protocol • Device inbound interface • Device outbound interface • Additional data (if any) • Base event IDs (if any) Field-based aggregation implements a more flexible aggregation mechanism; two events are aggregated if only the <i>selected</i> fields are the same for both events. Note: Do not include ESM <i>derived fields</i> in field-based aggregation. Note: Field-based aggregation creates a new alert that contains only the fields that were specified, so the rest of the fields are ignored, unless “Preserve Common Fields” is set to “Yes”. Field-based aggregation offers several advantages over basic aggregation, including: • Control over what fields to aggregate on • Start and end time set to the earliest start time and latest end time, respectively (instead of taking the values from the first event in the group, like basic aggregation) • Option to preserve common fields • Option to sum one or more numeric fields SmartConnector aggregation significantly reduces the amount of data received, and should be applied only when you use less than the total amount of information the event offers. For example, you could enable field-based aggregation to aggregate "accepts" and "rejects" in a firewall, but you should use it only if you are interested in the count of these events, instead of all the information provided by the firewall.

Configurable Settings, continued

Name Field	Value Field
Time Interval	Choose a time interval, if applicable, to use as a basis for aggregating the events the connector collects. Aggregation time interval and threshold settings need to both be set in order for the aggregation to be enabled. (Disabled , 1 sec, 5 sec, and so on, up to 1 hour.)
Event Threshold	Choose a number of events, if applicable, to use as a basis for aggregating the events the connector collects. This is the maximum count of events that can be aggregated; for example, if 150 events were found to be the same within the time interval selected (i.e., contained the same selected fields) and you select an event threshold of 100, you then receive two events, one of count 100 and another of count 50. This option is exclusive of Time Interval. (Disabled , 10 events, 50 events, and so on, up to 10,000 events.)
Field Names	Choose one or more fields, if applicable, to use as the basis for aggregating the events the connector collects. Use Ctrl+click to select multiple fields. The result is a comma-separated list of fields to monitor. For example, "eventName,deviceHostName" would aggregate events if they have the same event- and device-host names. You can use any of the event fields displayed in the event inspector; the name can contain no spaces and the first letter should not be capitalized.
Fields to Sum	Choose one or more fields, if applicable, to use as the basis for aggregating the events the connector collects. If specified, this set of numeric fields is summed rather than aggregated, preserved, or discarded. The most common fields to sum are bytesIn and bytesOut. Note that if any of the fields listed here are also in the list of field names to aggregate, they are aggregated and not summed.
Preserve Common Fields	(Yes No) Choosing Yes adds fields to the aggregated event if they have the same values for each event. Choosing No , the default, ignores non-aggregated fields in aggregated events.
Filter Aggregation	Filter Aggregation is a way of capturing aggregated event data from events that would otherwise be discarded due to an agent filter. Only events that would be filtered out are considered for filter aggregation (unlike Field-based aggregation, which looks at all events).
Time Interval	Choose a time interval, if applicable, to use as a basis for aggregating the events the connector collects. It is exclusive of Event Threshold. (Disabled , 1 sec, 5 sec, and so on, up to 1 hour.)
Event Threshold	Choose a number of events, if applicable, to use as a basis for aggregating the events the connector collects. This is the maximum count of events that can be aggregated; for example, if 150 events were found to be the same within the time interval selected (that is, contained the same selected fields) and you select an event threshold of 100, you then receive two events, one of count 100 and another of count 50. This option is exclusive of Time Interval. (Disabled , 10 events, 50 events, and so on, up to 10,000 events.)
Fields to Sum	(Optional) Choose one or more fields, if applicable, to use as the basis for aggregating the events the connector collects.
Processing	

Configurable Settings, continued

Name Field	Value Field
Preserve Raw Event	(Yes No) Some devices contain a raw event that can be captured as part of the generated alert. If that is not the case, most connectors can also produce a serialized version of the data stream that was parsed/processed to generate the ArcSight event. This feature allows the connector to preserve this serialized "raw event" as a field in the event inspector. This feature is disabled, by default, since using raw data increases the event size and therefore requires more database storage space. You can enable this by changing the Preserve Raw Event setting. If you choose Yes, the serialized representation of the "Raw Event" is sent to the selected destination and preserved in the Raw Event field.
Turbo Mode	If your configuration, reporting, and analytic usage permits, you can greatly accelerate the transfer of a sensor's event information through SmartConnectors by choosing one of two "turbo" (narrower data bandwidth) modes. Complete is the default transfer mode, which passes all the data arriving from the device, including any additional data (custom, or vendor-specific). This corresponds to <code>turbo.enabled=false</code> on the Manger. Since this value is not the default, be sure to add this property to the Manager's <code>server.properties</code> file. The first level of Turbo acceleration is called Faster and drops just additional data, while retaining all other information. The Fastest mode eliminates all but a core set of event attributes, in order to achieve the best throughput. Consider the possible effects such a restricted data set might have from a given device (for example, on reports, rules, threat resolution) before selecting it. The specific event attributes that apply to these modes in your enterprise are defined in the self-documented <code>\$ARCSIGHT_HOME/config/connector/agent.properties</code> file for the ArcSight Manager. Because these properties may have been adjusted for your needs, you should refer to this file for definitive lists. Only scanner SmartConnectors must run in Complete mode, to capture the additional data. Note: SmartConnector Turbo Modes are superseded by the Turbo Mode in use by the ArcSight Managers processing their events. For example, a Manager set to Faster cannot pass all the data possible for a SmartConnector that is set for the default of Complete.

Configurable Settings, continued

Name Field	Value Field
Enable Aggregation (in secs)	Note: If you have already used this feature for setting up previous SmartConnectors, you can continue to do so. However, ArcSight recommends that you use the new "Field Based Aggregation" on page 183 feature as a more flexible option. Here is the description of the legacy "Enable Aggregation" feature, for those who are still using it: When enabled, Enable Aggregation (in seconds) aggregates two or more events on the basis of the selected time value. (Disabled, 1, 2, 3, 4, 5, 10, 30, 60) The aggregation is performed on one or more matches for a fixed subset of fields: • Agent ID • Name • Device event category • Agent severity • Destination address • Destination user ID • Destination port • Request URL • Source address • Source user ID • Source port • Destination process name • Transport protocol • Application protocol • Device inbound interface • Device outbound interface • Additional data (if any) • Base event IDs (if any) The aggregated event shows the event count (how many events were aggregated into the displayed event) and event type. The rest of the fields in the aggregated event take the values of the first event in the set of aggregated events.
Limit Event Processing Rate	You can moderate the SmartConnector's burden on the CPU by reducing its processing rate. This can also be a means of dealing with the effects of event bursts. The choices range from Disabled (no limitation on CPU demand) to 1 eps (pass just one event per second, making the smallest demand on the CPU). Be sure to note that this option's effect varies with the category of SmartConnector in use, as described in the SmartConnector Processing Categories table below.
Fields to Obfuscate	Using MD5 hashing, this option allows you to specify a list of fields for obfuscation in a security event.

Configurable Settings, continued

Name Field	Value Field
Store Original Time In	This parameter allows you to move the original device receipt time to a specified field if altered by the time correction.
Enable Port-Service Mapping	If Enabled and one of the two fields destination port and application protocol is set, and the other is not, the one that is set is used to set the other. For example, if the destination port is 22 and application protocol is not set, then the application protocol is set to ssh. Default is Disabled.
Uppercase User Names	(Disabled Enabled) Default is Disabled. If set to any of the <i>enabled</i> settings, the two user name fields are automatically changed to uppercase. The original values are saved as follows: • Enabled (orig to ID) saves the original values to the sourceUserID and destinationUserID fields, respectively, overwriting any values that may have been there previously. • Enabled (orig to ID or Flex) saves the original values in the same fields if they do not already contain values, or to the flexString1 (source) and flexString2 (destination) fields if the ID fields do contain values. • Enabled (orig to Add. Data) saves the original values to additional data fields called OrigSrcUserName and OrigDstUserName, respectively. Note: The uppercase operation is typically done using the default Locale for the chosen platform. You can set this to a particular Locale by setting the connector.uppercase.user.name.locale property in agent.properties to the desired Locale (using "en_US" for U.S. English, for example).
Enable User Name Splitting	(Yes No) If this is set to yes and the destination user name contains commas in the event, this parameter duplicates that event. Each user name in the list is placed in one of the events. For example, if the destination user name in an event is "User 123, User 456", then that event is sent twice, with the destination user name set to "User 123" in the first and "User 456" in the second. Default is No.
Split File Name into Path and Name	(Yes No) If this is set to yes and an event's file name field is set but its file path field is not, this parameter splits the file name into a path and a name, placing each part into appropriate fields. For example, if the file name field is set to C:\dir\file.ext and the file path is not set, then the file path is set to C:\dir and the file name to file.ext. The separator character can be either \ or / as the system looks to the SmartConnector to determine its platform. Default is No.

Configurable Settings, continued

Name Field	Value Field
Event Integrity Algorithm	(Disabled SHA-256 SHA-1 MD5 SHA-512) If this is set to one of the algorithms (such as SHA-256), <i>and</i> the Preserve Raw Event parameter is Enabled, then additional event integrity internal events are generated, normally at a rate of about 1 per 50 normal events. The crypto signature field is <i>also</i> set in each event in the format: "<i>#seq(alg):digest</i>", where <i>seq</i> is a persistent event sequence number, <i>alg</i> is the message digest algorithm, and <i>digest</i> is the hexadecimal message digest. These extra events and the crypto signature field values can be used to verify that no events were tampered with after generation. Supported algorithms are: SHA-256, SHA-1, MD5, and SHA-512. Default is Disabled (that is, no algorithm is applied).
Generate Unparsed Events	(Yes No) If set to yes and some incoming event data cannot be parsed (perhaps because a device has been upgraded since the SmartConnector parser was written), then a special event named "Unparsed Event" is generated. The raw event appears in the event message field. If set to No, the SmartConnector log files indicate the unparsed events. Default is No.
Preserve System Health Events	(Yes No) If set to yes, internal system health events are preserved. SmartConnectors generate system health events that provide information about the systems on which they are installed (for example, disk usage, network memory, JVM memory, percentage of processing of CPU memory usage, and so forth). By default, these events are not retained or passed on to ArcSight destinations and, therefore, not available for viewing. Setting this option to yes makes them available in the Console.

Configurable Settings, continued

Name Field	Value Field
Enable Device Status Monitoring (in millisec)	(<NumberOfMilliseconds> -1 (disabled)) If set to a <NumberOfMilliseconds>, the selected SmartConnector generates internal events periodically 1 minute (60000 milliseconds) or greater with the status of the devices for which the connector is receiving normal events. These events have the name "Connector Device Status." Enabling periodic device status monitoring events helps monitor both the SmartConnector and device uptime. Device status monitoring events include this information, if available: • Event name (Connector Device Status) • Vendor and Product information • Source Address and Host Name • Zone • Last event received • Total number of events for the device since the connector started • Event count since last call Device status monitoring events can be set to generate every 1 minute (60000 milliseconds), or less frequently (that is, a greater number of milliseconds than the minimum). If you specify less than 60000, you get a warning in the log that the minimum is 60000 milliseconds (1 minute) and the system uses the minimum. If you enter a non-number in the field, it generates an error in the log that the value could not be parsed. In this case, the feature is disabled (and logged as such). In such cases, there is no indication on the Console that anything went wrong because there is no way for the Connector to convey that error.
Payload Sampling (when available)	Some SmartConnectors use Payload sampling to send a portion of packet payload (as opposed to the complete payload) along with the original event. This portion is retrieved using the on-demand payload retrieval in the event inspector.
Maximum Length	You can configure the maximum length of the payload sample using the following values: • Discard • 128 bytes • 256 bytes • 512 bytes • 1 Kbyte When the Discard option is chosen, no payload sample is sent inside the original event.
Mask Non-printable Characters	This feature allows you to mask the non-printable characters in the payload sample.

Configurable Settings, continued

Name Field	Value Field
Filters	Agent severity is the translation of the device severity into normalized values. For example, some connectors use a device severity scale of 1-10, whereas others use a scale of high, medium and low. These values are normalized into a single agent severity scale. The default scale is Low, Medium, High, and Very High. An event can also be classified as Unknown if the data source did not provide a severity rating.
Filter Out	Filters for SmartConnectors are exclusive (filter out). Events that meet the connector filtering criteria are not forwarded to the destination. During SmartConnector set up, you can configure the connector to use filter conditions that do not pass events to the destination according to specific criteria. For example, you can use filters to exclude events with certain characteristics or events from specific network devices.
Very High Severity Event Definition	Enter a filter condition to sort for very high severity events.
High Severity Event Definition	Enter a filter condition to sort for high severity events.
Medium Severity Event Definition	Enter a filter condition to sort for medium severity events.
Low Severity Event Definition	Enter a filter condition to sort for low severity events.
Unknown Severity Event Definition	Enter a filter condition to sort for unknown severity events.

SmartConnector Processing Categories

The following table shows the SmartConnector processing categories.

SmartConnector Processing Categories


SmartConnector Type	Effects of Limited Usage
Syslog connectors	Due to the nature of UDP (user datagram protocol, the transport protocol used by Syslog), these connectors can potentially lose events if the configurable event rate is exceeded. This is because the connector delays processing to match the event rate configured, and while in this state, the UDP cache may fill and the operating system drop UDP messages.
SNMP connectors	Similar to Syslog connectors, when the event rate is limited on SNMP connectors, they potentially lose events. SNMP is also UDP-based and has the same issues as Syslog.

SmartConnector Processing Categories, continued

SmartConnector Type	Effects of Limited Usage
Database connectors	Since connectors "follow" the database tables, limiting the event rate for database connectors can slow the operation of other connectors. The result can be an event backlog sufficient to delay the reporting of alerts by as much as minutes or hours. However, no events are lost, unless the database tables are truncated. After the event burst is over, the connector may eventually catch up with the database if the event rate does not exceed the configured limit.
File connectors	Similar to database connectors, file-based connectors "follow" files, so limiting their event rates also causes an event backlog. This can eventually force the connector to fall behind by as much as minutes or hours, depending on the actual event rate. Similarly, the connectors may catch up if the event rate does not exceed the configured rate.
Proprietary API connectors	These connectors' behavior depends on the particular API, (for example, OPSEC behaves differently than PostOffice and RDEP). But in most cases, there is no event loss unless the internal buffers and queues of the API implementation fill up. Therefore, these connectors work much like database or file connectors.

SmartConnector Time Interval Options

This time interval applies to the Alternate Settings and it specifies when the alternate settings must be used by the SmartConnector. For example, if you want to cache the events during the day and send everything at night, you can configure the Transport Mode to cache in the default configuration and configure the Transport Mode to normal in the Alternate Settings, then you would set the time interval from 8PM to 8AM (next day).



- **From:** Specifies the starting time to apply the Alternate settings.
- **To:** Specifies the ending time for the Alternate settings, when it reverts to the default settings). If this is less than the From setting, the value is interpreted as the next day. For example, a setting from 8PM to 8AM is interpreted as starting at 8PM and ending at 8AM the following day.

To save configuration changes to the SmartConnector, click **OK**.

Managing SmartConnector Filter Conditions

SmartConnectors have built-in filters that function as a filtering tool between devices and the ArcSight Manager, using filtering conditions. Filtering conditions are set with a combination of AND or OR statements and data field values. Extra or unneeded events are filtered out to minimize the number of events sent to the ArcSight Manager and analyzed in the ArcSight Console.


Adding SmartConnector Filter Conditions

1. In the Navigator panel, choose the **Connectors** resource tree.
2. In the Connectors resource tree, right-click a SmartConnector and select **Configure**.
3. Go to the **Default > Filter** tab.

The Filters tab displays a lists the default filters for the SmartConnector.

4. Select a filter for which you want to customize a condition.

The Common Conditions Editor appears.



5. Right-click the event alias, **event1**, and select **New Condition**.
6. In the popup, select a data field from the drop-down menu and start defining your statement. For help with defining condition statements, see ["Common Conditions Editor \(CCE\)" on page 887](#), especially ["Condition Tree Command Buttons" on page 890](#) and ["Condition Tree Context Menu Commands" on page 892.](#))

7. Select a logical operator on the drop-down menu. See ["Logical Operators" on page 1006](#) for details.
8. Enter a value in the last text field to complete the conditional statement.
9. Click **OK**.

Deleting SmartConnector Filter Conditions


1. In the Navigator panel, choose the **Connectors** resource tree.
2. In the Connectors resource tree, right-click the SmartConnector and choose **Configure**.
3. In the Filtering section on the Advanced tab, right-click a condition and choose **Delete condition**.

Setting Special Severity Levels

You can customize or conditionalize the event-severity levels reported by SmartConnectors. Customizing means pre-setting a given SmartConnector's filter to one specific severity level; conditionalizing is essentially the same, but with the addition of a filter condition to determine when the pre-set severity level is reported.

To configure a custom or conditional severity level:

1. Choose the **Connectors** resource tree in the Navigator panel.
2. In the Connectors resource tree, right-click the appropriate SmartConnector and choose **Configure**.
3. In the Connector Configuration Editor, select the **Connector: Default: Filters** tab.
4. In the Filters tab, select a severity level.
5. In the Filter Condition dialog box choose a field, a logical operator, and enter a value for the condition.



- Click **OK** in the Filter Condition dialog box and **Apply** or **OK** in the Connector Configuration Editor.

In the example, we selected the **"Very-High Severity"** filter and defined a condition in which Category Significance contains **Hostile**. When this condition is met, the severity of the event becomes "Very-High."

For more information, see ["Managing SmartConnector Filter Conditions" on page 192](#).

Sending Model Mappings to SmartConnectors

Updates to network model mappings are sent automatically from the ArcSight Manager to SmartConnectors within heartbeat messages. The heartbeat messages themselves are sent on an interval which can be anywhere from every 5 seconds to every 10 minutes, but network model mappings are included in the messages only when there are updates to the model.



Note: The interval on which information is exchanged between the Manager and SmartConnectors is determined by the Heartbeat Frequency setting on each Connector. (See information on ["Heartbeat Frequency" on page 180](#) in default content tab configuration fields under ["Selecting and Setting SmartConnector Parameters" on page 175](#).)

If you have made several configuration updates to the network model on the Manager and would like these changes to take effect immediately on the SmartConnectors without waiting for the next automatic refresh, you can use the following command to send the update information to a selected Connector.

Where: Navigator > Resources > Connectors

To send model mappings to SmartConnectors:

Right-click the SmartConnector you want to update and choose **Send Model mappings now**.

This sends information about the current network model mappings from the Manager to the selected Connector. It forces a comprehensive refresh of the zone mappings and network model information on the Connector.

Sending Control Commands to SmartConnectors

From the Console, you can issue basic event-flow-control commands to SmartConnectors, get the operational status of a SmartConnector, or issue control commands to network devices through their SmartConnectors. This topic discusses the first two points.

To author rule-driven device-command responses to events, see ["Adding, Editing, or Removing a Rule Action" on page 533](#).

Getting Connector Status

You can see a SmartConnector's current operational state at any time.

Where: Navigator > Resources > Connectors

To get a connector status report:

1. Right-click the SmartConnector.
2. Select **Send Command>Status>Get Status**.

The Connector Status window displays all the connector's current parameters. You have the option to save the information in a text file by clicking the **Save** icon.

Sending Standard Flow-Control Commands



Note: Familiarize yourself with the effects of certain commands described in this topic.

- Because there is no local cache, events that occur while a connector is stopped or paused are not retained.
- If a SmartConnector runs out of disk space, it can lose its ability to track events.
- The **Terminate** command should only be used in very special circumstances because it **kills all** SmartConnector processes.
- See ["Adding, Editing, or Removing a Rule Action" on page 533](#) for a description of the rule-based automated alternative for giving SmartConnector commands.

Where: Navigator > Resources > Connectors

1. Right-click the SmartConnector, choose **Send Command**, and one of the following menu options described below.



Note: Commands available on this menu vary depending on which SmartConnectors you are using. The standard set of commands is described here.

SmartConnector Status

Category	Command Descriptions
Status Category	Get Status Provides a full report on the selected SmartConnector's current operational state.
	Get Device Status Provides the status of the device that reports to the SmartConnector. (Currently only available for the CiscoIDS/IPS SmartConnector.)
Connector Process Category	Restart Stops, then restarts a running SmartConnector. Caution: After a connector is terminated, Console commands cannot access it. Therefore, a restart works only on a connector that is currently running.
	Terminate Shuts down the SmartConnector and all processes the SmartConnector started. Caution: Once a connector is terminated, Console commands (including Connector Process > Restart) cannot access it. The connector must be restarted manually from the machine on which it is installed.
	Event Flow Category
Event Flow Category	Pause Stops the SmartConnector from sending events to the ArcSight Manager. Note: Events received from the target device are saved in the connector cache (even though the connector is in Pause state).
	Stop Stops the SmartConnector from sending events to the ArcSight Manager. Caution: A Stop command causes the SmartConnector to stop sending all events, which means it is not sending events to the connector cache, either.
	Start Prompts the SmartConnector (previously in Stop or Pause state) to start sending events to the ArcSight Manager.
	Network Category
Network Category	Flush Name Resolver Cache Clears cache for Network name resolver.

SmartConnector Status, continued

Category	Command Descriptions
Upgrade Category	Upgrade Launches a Command Parameters dialog for remote upgrade to newer versions of ArcSightSmartConnectors for managed assets. Provide the version number of the connector to which you want to upgrade and a wait time to verify that the upgrade completed successfully. (If the upgrade is not successful, the system performs an automatic rollback to the previous version of the connector.) Click OK to start the upgrade. See "Upgrading SmartConnectors" on page 209 for prerequisites for the upgrade process and detailed information on how to upgrade Connectors.
	Rollback Upgrade Launches a Command Parameters dialog for remote rollback of connector version to a specified previous version. See "Upgrading SmartConnectors" on page 209 for complete information.
Adjust Category	Rename Mismatched Override Files Enables you to remotely rename an connector parser override file whose version stamp no longer matches the parser that it was intended to override. Renaming it appends ".1" (or 2, or 3, if earlier numbers are in use), which stops the file from being used. The first parameter is a regular expression you can specify to match specific override files (or blank, the default, for all). The second parameter is a boolean where true, the default, means restart the connector if any files are renamed.

The Console's status bar shows a confirmation message when the flow control option takes effect.

Tech Support Commands

Tech Support commands are provided for use primarily by ArcSight Customer Support. Brief descriptions of these Tech Support commands are provided for informational purposes, but these commands are not intended for use by ArcSight customers except as instructed by Customer Support.

Tech Support Category Commands

Command	Description
Get Support Info	Gets logs and other feedback on connectors.
Get 'agent.properties'	Shows the list of properties for the selected connector.
Get Upgrade Logs	Gets upgrade logs on connectors.
Get 'agent.wrapper.conf'	Shows the wrapper configuration for the selected connector.
Get Configuration XML File	Shows the XML configuration file for the selected connector.
Get Thread Dump	Gets one thread dump for the selected connector.

Tech Support Category Commands, continued

Command	Description
Get Two Thread Dumps	Gets two thread dumps for the selected connector spaced by the time interval specified. By comparing both thread dumps, Customer Support can troubleshoot connectors with threads that are hanging for unknown reasons.
Get Heap Dump	This generates a heap dump, if possible, which can be useful to ArcSight in some situations to analyze problems. The destination ID is used as part of the file name, the file is placed in the same directory as the connector's logs, and normally only 10 such files are kept.
Get last N lines of 'agent.log'	Shows an excerpt from the connector log file based on the number of lines you specify. The default is 500 lines.
Get System Properties	Shows system properties for the selected connector, including details on variables such as Java runtime name, Java virtual machine (VM) version, operating system name, paths for various Java components, paths for ArcSight Home, user directories, user home, and so forth.
Enable Event Flow Tracing	Allows you to specify a component and fields to log for initiating an event flow trace. Component and field names must be provided per appropriate syntax. The component should be chosen from the components listed in the Get Status results.
Disable Event Flow Tracing	Disables event flow tracing on the selected component.
Get Event Flow Tracing Log	When tracing is enabled on the selected connector, the connector logs data about events it receives.
DNS Test	This command takes one parameter, which is either a host name to resolve or an IP address to reverse resolve. This is useful to see what results would normally be expected for the name resolver component of the connector, since it uses the same mechanism to do the lookup as the name resolver uses.
Enable Map File Logging	Directs the AgentNATProcessor component, which processes map files for each event, to log what it is doing for each event. By default the last 100 events are logged.
Disable Map File Logging	Directs the AgentNATProcessor to stop logging.
Get Collected Map File Logging	Gets the collected log messages for the most recent events (100, by default), which may help debug problems with why a map file is not operating as expected.


Mapping Commands for Additional Data Fields

The following commands provide access to SmartConnector component mapping and event categorization for advanced users.


Mapping Commands

Command	Description
Mapping Category	
Get Additional Data Names	Returns a list of additional data names seen for each device vendor/product combination since the connector started running. For example: <pre>Additional Data Names Seen: Generic (no vendor/product): test1 [3 times] test11 test13 [2 times] Vendor/product [vend/prod]: test1 test10 [6 times]</pre> By default, the command limits the list to show only the most recent 100 device vendor/product combinations and the most recent 100 names for each. Tip: You can change this limit by editing the SmartConnector property <code>agent.additionaldata.mapper.track.max.names</code> in the file <code>\$ARCSIGHT_HOME/ArcSightSmartAgents/current/user/agent/agent.properties</code> on the machine where the connector is installed. In most cases we recommend keeping the defaults. If you change a property setting such as this, restart the connector. If a data name is not a string, its data type is displayed in the list. If the connector saw an additional data name more than once, the command output indicates the number of times the name was seen.

Mapping Commands, continued

Command	Description
Map Additional Data Name	Brings up a dialog where you can map an additional data name for the selected connector. If you are using additional data, add the <code>turbo.enabled=false</code> property to the Manager's <code>server.properties</code> file. If you are unfamiliar with properties file setting, see "Managing and Changing Properties File Settings" in the ESM Administrator's Guide.  For a generic mapping, you can leave the Device vendor and Device product fields blank. For a specific mapping, fill in these fields with the appropriate vendor and product names. Typically, the Additional data name is one of the names shown in the Get Additional Data Names output (but can be another name not on that list). The ArcSight field must be a valid ArcSight event field. Click OK to create the mapping. Here is an example of the command output for a successful generic mapping: <pre>Successfully mapped additional data name [test11] to event field [message] for vendor/product []</pre> A successful device vendor/product-specific mapping returns output similar to this: <pre>Successfully mapped additional data name [test10] to event field [message] for vendor/product [vend/prod]</pre> If the additional data name has not been seen, the name is still mapped, but with a warning like this: <pre>Successfully mapped additional data name [foo] to event field [deviceCustomString1] for vendor/product [vend/prod] (note that additional data name [foo] has not been seen for vendor/product [vend/prod])</pre> If the ArcSight field is not valid, the error returned is similar to this: <pre>Failed to map additional data name [bar] to event field [messages] for vendor/product [vend/prod] (event field [messages] is unknown)</pre>

Mapping Commands, continued

Command	Description
Unmap Additional Data Name	Brings up a dialog where you can unmap an additional data name for the selected connector.  To remove a generic mapping, you can leave the Device vendor and Device product fields blank. To remove a specific mapping, fill in these fields with the appropriate vendor and product names. The additional data name should be one that was previously mapped for the specified device vendor and product combination. Click OK to unmap the data name. Here is an example of the command output for a successful generic unmapping: <pre>Successfully unmapped additional data name [test11] for vendor/product []</pre> A successful device vendor/product-specific unmapping returns output similar to this: <pre>Successfully unmapped additional data name [foo] for vendor/product [vend/prod]</pre> If the specified additional data name was not previously mapped, the output looks like this: <pre>Failed to unmap additional data name [foo] for vendor/product [vend/prod] (not previously mapped)</pre> Notes: • One additional data name can be mapped to more than one ArcSight field for the same device vendor/product combination, and in this case unmapping it unmaps it from all ArcSight fields for that device vendor/product. This is an unlikely scenario, however. • The converse case, where multiple additional data names are mapped to the same ArcSight field for the same device vendor/product combination, results in the last mapping taking precedence over any previous mappings to that ArcSight field for that device vendor/product. No warning is generated in this case.
Categorizer mapper Category	

Mapping Commands, continued

Command	Description
Reload custom categorizations	There are several ways to set event category information for events. The least common of these is to store custom categorization files (organized by vendor and product) on the connector machine in the <code>user/agent/aup/acp/categorizer/current</code> directory (or the <code>user/agent/acp/categorizer/current</code> directory). If such categorization files exist and have been changed, this command reloads them without restarting the connector.
Reload custom map files	Rescans and reloads map files in the <code>user/agent/map</code> directory on the machine where the connector is installed. The map files are named in the form <code>map.n.properties</code>, where <i>n</i> is a number starting with 0. Use this command to immediately apply the latest changes. Not all connector setups include custom map files. Caution: Map files are created on some connector machines to fulfill specific needs. If you are not familiar with the categorizer/mapping setup of an environment, we recommend that you do not use these commands.
Reload external map files	Re-scans and reloads external map files in the <code>user/agent/extmap</code> directory on the machine where the connector is installed. The map files are named in the form <code>extmap.n.properties</code>, where <i>n</i> is a number starting with 0. Use this command to immediately apply the latest changes. Not all connector setups include custom external map files. Caution: External map files are created on some connector machines to fulfill specific needs. If you are not familiar with them, we recommend that you do not use Reload commands.


Managing SmartConnector Groups

You can best manage ArcSight SmartConnectors when you organize them into groups. You'll find all uncategorized SmartConnectors in the Unassigned group.

You can move or copy groups and SmartConnectors into other groups in the Connectors resource tree by using drag-and-drop. If a group is deleted, the SmartConnectors within that group are also deleted.

You should not delete a Connector resource at the ArcSight Console, unless the corresponding SmartConnector is first stopped. If the SmartConnector on the device is running and its Connector resource is deleted, the SmartConnector cannot send events to the ArcSight Manager, causing the SmartConnector to start caching events and eventually dropping these events.

Where: Navigator > Resources > Connectors



To create a SmartConnector group:

1. Right-click a group and choose **New Group**.
A "name" text field appears under the group you selected.
2. In the "name" text field, type in a name.
3. Press **Enter**.

To rename a SmartConnector group:

1. Right-click a group and choose **Rename**.
2. In the "name" text field, rename the group.
3. Press **Enter**.

To edit a SmartConnector group:

1. Right-click a group and choose **Edit Group**.
2. In the Group Editor, edit the **Name** and **Description** text field.
3. Click **OK**.

To move or copy a SmartConnector group:

1. Navigate to a group and drag and drop it into another group.
2. Choose **Move** to move the group, **Copy** to make a separate copy of the group, or **Link** to create a copy of the group that is linked to the original group.

If you choose **Copy**, you create a separate copy of the group that is not affected when the original group is edited. If you choose **Link**, you create a copy of the group that is linked to the original group. Therefore, if you edit a linked group, whether it be the original or the copy, all links are edited as well. When deleting linked groups, you can either delete the selected group or all linked groups.

To delete a SmartConnector group:

1. Right-click a group and choose **Delete Group**.
2. In the dialog box, click **Yes**.

Importing and Exporting SmartConnector Configurations

You can import and export SmartConnector configurations as a means of sharing custom configurations among several connectors on the same or multiple Managers. Rather than redefining a complex configuration on each connector, you can export the configuration as an XML file and then import it into connectors that share some or all of its configuration settings. An override feature allows you to make changes to any of the parameter values upon import.

Importing a SmartConnector Configuration

 **Tip:** Save and import SmartConnector configurations as XML files.

Where: Navigator > Resources > Connectors

1. Right-click the SmartConnector into which you want to import a new configuration and choose **Import Connector Configuration**.
2. In the file browser, navigate to and select the .xml file that contains the connector configuration, and click **Open**.

This displays a dialog showing original and proposed new configuration settings for the selected connector, with an option to override any of the proposed new values.

3. Click **Show** to show the details of the import or **Hide** to hide them.
 - a. Limit the import to only filters by keeping all values in the original configuration and choosing to override only the filter values with the imported values as is detailed in SmartConnector Filters.

Scroll down to the Filters section at the end of the Import dialog to see the filters.

When you are satisfied with the settings to import and overrides (if any), click **OK** to import the configuration.
 - b. This dialog shows original values for the selected connector configuration and new values applied upon import. You can override any of the settings you do not want to import by either keeping the parameter value in the original configuration or defining a new value.
 - c. On the Import Connector Configuration dialog, review the import information and override any values that you do not want to import.

Exporting a SmartConnector Configuration



Tip: SmartConnector configurations must be saved as XML files.

Where: Navigator > Resources > Connectors

1. Right-click the SmartConnector you want to export, and choose **Export Connector Configuration As**.

This brings up a file browser where you can navigate to the location where you want to save the configuration as an XML file.

2. In the file browser, navigate to and select the location where you want to save the configuration, provide a name for the file, and click **Save**.

SmartConnector Filters

You can import and export only the filters associated with SmartConnectors as a part of an import or export on a SmartConnector.

To export a SmartConnector filter, export the connector that uses the filter as described in ["Exporting a SmartConnector Configuration" above](#).

To import a SmartConnector filter into another connector:

1. Select the the SmartConnector to which you want to add a new filter.
2. Follow the steps in ["Importing a SmartConnector Configuration" on the previous page](#).
3. On the Import Connector Configuration dialog, scroll down to the Filters section at the end of the Import dialog to see the SmartConnector Filters.
4. Limit the import to only the filters you want by keeping all values in the original configuration and choosing to override only the filter values with the import.
5. Click **OK**.

Using Additional Data Fields

Some devices include event data with data fields that are not included in the standard event data schema. You can configure your SmartConnector to send these fields as additional data and map them to ESM schema fields. The mapping can vary based on the device vendor and product and can be controlled from the ArcSight Console, with the mappings stored on the SmartConnector machine.


Prerequisite:

Add the `turbo.enabled=false` property to the Manager's `server.properties` file. If you are unfamiliar with properties file setting, refer to the *ESM Administrator's Guide*, topic on "Managing and Changing Properties File Settings."


To map additional data fields:

Where: Navigator > Resources > Connectors

1. Expand a connector group and locate the SmartConnector you want to configure for additional data field mapping. Make sure this connector is in the *running* state.
2. Right-click the connector and select **Send Command > Mapping > Get Additional Data Names**.



3. See [Get Additional Data Names](#) sample output. Take note of the field names of interest. You will enter these fields for mapping later in this procedure.
4. Right-click the connector again and select **Send Command > Mapping > Map Additional Data Name**. The following is displayed:

A screenshot of a dialog box titled 'Map Additional Data Name'. It contains a table with two columns: 'Name' and 'Value'. The 'Name' column has a dropdown menu with 'Misc' selected. The 'Value' column has four rows for 'Device vendor', 'Device product', 'Additional data name', and 'ArcSight field'. At the bottom right of the dialog are 'OK' and 'Cancel' buttons.

5. Enter values, for example:

Field	Example value
Device vendor	For example, Cisco
Device product	For example, Cisco Secure ACS
Additional data name	Field name from the connector to be mapped. Taken from the generated name list from the previous step
ArcSight field	ArcSight field to store the additional data

See also [Map Additional Data Name](#) for additional instructions on entering values in the mapping dialog, additional instructions for changing the `server.properties` file, and sample messages.

6. To verify the mapping, right-click the connector again and select **Send Command > Status > Get Status**. Check the `NGCustomAdditionalDataMapper<n>` variables in the output. For example:

```
NGCustomAdditionalDataMapper0.....Generic
mappings:test11=>message
NGCustomAdditionalDataMapper1.....Mappings for
vend/prod:test10=>message, foo=>deviceCustomString1
```



Note: Only mappings for loaded device vendor/product combinations are included. This includes mappings for vendor/product combinations that have had mapping or unmapping commands executed (even unsuccessful ones), and vendor/product combinations for which additional data-laden events have been seen. Unloaded mappings on disk are not included.

Create a map file on the SmartConnector host:

Connector map files are used to map a field value to another field, and more. For proper map file construction, follow the information in the *FlexConnector Developer's Guide*, which is available from [the Open Text Community](#). Refer to the topic, "Map Files."

Map files follow the filename format, `map.X.properties`, where **X** is the next sequential number following any other existing map file in that directory.


Store the map file in the SmartConnector host in this directory:

```
$ARCSIGHT_HOME/user/agent/map/map.X.properties
```

Updating this file does not require any connector restarts.

To display additional data fields in the active channel:

1. Create an active channel and specify the event fields to be displayed. See "[Creating or Editing an Active Channel](#)" on page 224 for instructions.
2. Right-click on the column header and select **Columns > Add/Remove Column > Additional Data**.



3. Select the additional data field columns to be added to the active channel.

Once the additional data fields are available on the channel, you can select the event (the row) and look at the Event Inspector panel for event details. Scroll all the way down the event details and look for the **Additional Data** category.



Note: Limitations

- Additional data fields cannot be part of a field set.
- Additional data fields cannot be viewed on the ArcSight Command Center.

Upgrading SmartConnectors

You can centrally manage and configure SmartConnectors and update them remotely. Use the **Upgrade** command on the ArcSight Console to upgrade to newer versions of SmartConnector software for managed devices. Use the **Rollback** command to revert to a previous version on an upgraded connector.

The Upgrade command lets you launch, manage, and review the status of upgrades for all SmartConnectors. A fail-over mechanism launches SmartConnectors with previous versions if

upgrades fail. All communication and upgrade processes between components (Console, Manager, connectors) take place over secure connections.

The ArcSight Console reflects current version information for all of your SmartConnectors.



Note: SmartConnector remote upgrade is supported for connectors installed on Linux and Windows platforms only.


Overview of the Upgrade Process

1. As an ArcSight customer, you receive e-mail notifications about new connector releases from Customer Support.
2. ArcSight administrators download the latest releases to the ArcSight Manager where they are available for SmartConnector upgrades.



Tip: SmartConnector upgrade version files are delivered as ArcSight Update packs (.aup) files. (ArcSight update packs are compressed file sets, similar to .zips.) The administrator copies the .aup file to ARCSIGHT_HOME/updates/ onto a running ArcSight Manager. The Manager automatically unzips the .aup file and copies its contents to ARCSIGHT_HOME/repository/.

3. From the ArcSight Console, administrators select connectors to be upgraded (one at a time) and launch the upgrade command for each of them.
4. The selected connectors upgrade themselves, restart, and send upgrade results (success or failure) back to the ArcSight Console through the ArcSight Manager. The upgraded connector runs in the same home directory as the old connector.
 - If the upgrade is successful, the new connector starts and reports on successful upgrade status.
 - If the upgraded connector fails to start, the original connector restarts automatically as a fail-over measure. This is essentially an automatic rollback, and restart.





Tip: Monitoring SmartConnector Upgrade Status

SmartConnectors automatically determine their upgrade status when they start.

- When a connector starts up, it determines whether it is upgraded.
- If so, it waits for a configurable time interval for events from the monitored device to be processed.
- If, after that time interval, events have been processed, the SmartConnector is deemed up and running. The ArcSight Console indicates that the upgrade for that connector is a success and the newer connector version is reflected.

Additional notes on the SmartConnector upgrade procedure:

- When upgrading SmartConnectors, be sure to download current versions of the connector configuration guides from the [ArcSight Product Documentation](#). New or revised information is provided in these guides as appropriate per each release of SmartConnectors. To check version numbers on your current connectors, see "[Getting Status and Versions on Installed SmartConnectors](#)" on page 213.
- You need administrative permissions to upgrade connectors.
- Newer versions of the connectors you want must be available on the Manager to which you are connected.
- The option for remote upgrade is available only on SmartConnectors of version 4.0.2.xxxx.0 or newer. Earlier versions of Connectors (or Agents) must be upgraded manually by installing a newer version of the connector.
- As a prerequisite to upgrading connectors, both the ArcSight Manager and the connector must be running.
- The Upgrade SmartConnectors command is available as one of several SmartConnector control commands.

SmartConnector Upgrade Procedure

Where: Navigator > Resources > Connectors

1. Select the connector you want to upgrade, right-click to bring up the context menu, and choose **Send Command > Upgrade > Upgrade**.

This opens a Command Parameters dialog.

2. Provide the following information in the dialog.
 - **Version** - The Version field provides a drop-down menu showing the connector versions available on this Manager. Choose the Version number of the connector to which you want to upgrade.

- **Event wait (sec)** - Number of seconds the upgrade process waits for the first event from the device after the new, upgraded connector is started. If no events are received from the device within the specified time frame, the upgrade is considered "failed" and the old connector is launched.

This optional check is an additional safeguard against upgrade failures. For example, the connector binaries may have been upgraded successfully, but the new version may have problems communicating with the device. In that case, this check assumes that the upgrade failed and bring back the old connector.

If the **Event wait (sec)** value is **0** (the default), then the upgrade does not perform this check.

3. Click **OK** to close the dialog and start the upgrade.

As the upgrade proceeds, the connector shows as down and then running again in the resource tree. Status messages on the ArcSight Console indicate whether the upgrade succeeds or fails. You can check the logs for the connector to determine if the upgrade succeeded. (**Send Command > Tech Support > Get 'agent.properties'** and **Get Upgrade Logs.**)

Rolling back to a Previous Version

- You need administrative permissions to roll back Connectors.
- Rollback automatically reinstates the most recent version prior to the currently installed version. You cannot do a remote rollback on a connector to other than the previously installed version.
- Check with OpenText Customer Support for ArcSight products on the SmartConnector version to which you can roll back.

Where: Navigator > Resources > Connectors

1. Select the connector you want to upgrade.
2. Right-click and select **Send Command > Upgrade > Rollback** from the context menu.

As the rollback proceeds, the connector shows as down and then running again in the resource tree. You can check the logs for the connector to determine if the rollback succeeded. (**Send Command > Tech Support > Get 'agent.properties'** and **Get Upgrade Logs.**)

Troubleshooting

If an upgrade or rollback fails, you can review the related logs. Choose **Send Command > Tech Support > Get Upgrade Logs** from the ArcSight Console menus.

You can also use the Send Logs wizard to collect and compress logs, which you can then attach to an email, if desired.

Getting Status and Versions on Installed SmartConnectors

Before or after you upgrade a SmartConnectors, you may want to check version numbers of currently installed connectors or get other status information. There are several ways to get information on currently installed connectors (including various control commands, channels, dashboards). Two of these are highlighted here as easy ways to get connector version information.

Where: Navigator > Resources > Connectors

To get status on a SmartConnector:

1. Select the connector you want to upgrade.
2. Right-click and select **Send Command > Status > Get Status** from the context menu.

The Status information on a connector includes "Agent Version" near the top of the message window. The information includes, but is not limited to, the content version, connector type, responses processed, and so on.

To view SmartConnector Dashboards:

1. Select **Dashboards** from the Navigator panel.
2. Expand the folders to find various dashboards. To view a dashboard, right-click it and choose **Show Dashboard**.

You can find SmartConnector dashboards in /Dashboards/Shared/All Dashboards/ArcSight Administration/Connectors/System Health.

Consuming Events from Transformation Hub

ESM can consume events from the ArcSight Transformation Hub. This is done first by configuring the SmartConnector as follows:

During SmartConnector setup:

1. Add the **ArcSight Manager** destination as described in the *ArcSight SmartConnectors User's Guide*.

Additionally, set these options:

Option	Setting	Purpose
AUP Master Destination	true	To enable event enrichment with zone information. If the setting is true, ESM sends zone information back to the connector. The connector then adds the zone information in events forwarded to ESM.
Filter Out All Events	true	To filter out and aggregate, and therefore reduce event volume sent to ESM. This increases efficiency and reduces processing time.

2. Save the configuration.
3. Add another destination. This time, select Transformation Hub, also described in the ArcSight SmartConnectors User's Guide.

For Content Format, select ESM Binary or Avro. This configures the connector to send events to Transformation Hub in binary format or Avro format.

During ArcSight Manager setup:

1. Make sure you configure ESM to connect to Transformation Hub.
2. Specify the same topic as in the SmartConnector.

If this configuration was not done during ESM installation, refer to the [ESM Administrator's Guide](#) for procedures to use the `managersetup` command and launch the configuration wizard again.

On the Console, the Transformation Hub connector is represented as follows:

Icon	Color	Description
	Green	Indicates that • The connector has been set up with ArcSight Manager as destination with AUP Master Destination = true, and Filter Out All Events = true. • The connector has been set up with Transformation Hub destination with "ESM Binary" or "Avro" as content format with topic also specified. • ArcSight Manager has been set up as an Transformation Hub consumer.
	Grey	Indicates that the Manager setup is <i>not</i> configured as a Transformation Hub consumer even if the SmartConnector itself is set with the appropriate Transformation Hub destination, content type, and topic. Therefore, Manager is not consuming events coming in from Transformation Hub.



Note: Transformation Hub configurations on the SmartConnector side *and* the Manager side *must* match.

If the SmartConnector is *not* set up with the Transformation Hub destination for ESM, even if the Manager is configured for Transformation Hub consumption during Manager setup, the icon on the Console will appear as a regular connector.

For more information about Transformation Hub, refer to the *ArcSight Open Data Platform Transformation Hub Guide*.

Chapter 7: Managing Notifications

Notifications and their content are created using rules configured with the Send Notification rule action (see ["Rule Actions Reference" on page 539](#)).

Managing Received Notifications

If the Notifications button in the ArcSight Console toolbar indicates that new notifications have arrived () you click that button to open the Notifications tab in the Viewer panel. This is your central notification repository if you belong to the destination group configured to receive notifications on the Console (the notification group's Destination Type is set to Console).

You can open the Notifications manager at any time by clicking the toolbar button, even if no new notifications are present.

To use the Notifications manager you first choose a category tab for the type of notification received.

Notification Category	Use
Pending 	These are notifications that you have not yet handled (reassigned to one of the following categories). Pending notifications older than 24 hours are automatically refiled as Not Acknowledged.
Undeliverable 	These are notifications that were not delivered.
Acknowledged 	These are notifications to which you have replied.
Not Acknowledged 	Pending notifications that go unacknowledged or unresolved for more than 24 hours are automatically refiled as Not Acknowledged.
Resolved 	These are notifications for which you or a colleague have found a resolution and so have marked the notification accordingly.
Informational 	These are notifications that are provided for information purposes only and do not require resolution or intervention. The Informational tab includes a Delete button. If you no longer need an informational notification, select it and click Delete.



Note: If you do not see notifications appearing, make sure your ESM user identity (not just your e-mail address) is set as a destination in the Notifications Editor.

In a category, click **Acknowledge** to mark a selected notification as acknowledged. Click **View Event** to see the event that triggered a notification. Click **Resolve** to reclassify the notification as Resolved.



Note: If a notification goes to multiple destinations, any user that belongs to those destinations can acknowledge or resolve it. The updated status is the same for all destinations.

For each category of notification there is a common set of columns of information concerning them.

Notification Column	Definition
Priority	This is the same priority set by the SmartConnector and modified by the current threat level formula (and seen in grid views), unless modified by the rule that triggered the notification.
Triggering Event	The event that caused the rule to trigger the notification.
Notification Group	The branch of the Notifications resource tree to which this destination belongs.
Escalation Level	The Escalation Level (and implied destinations) the notification has reached while waiting for resolution.
Create Time	The time at which the notification was created



Note: Also note that you can set a severity threshold for notification pop-ups and sounds in ArcSight Console Preferences.

Managing Notification Groups

On the Console, a notification group is denoted as a destination. This means this group is the recipient of notifications.

Purpose: The ArcSight Console has default destination groups that you can use. A destination group is assigned to users to receive specific notifications.

Where used:

- User attribute, described in ["Creating or Editing a User" on page 102](#)
- Rule action for sending notifications in case certain rule conditions are met. See ["Rule Actions Reference" on page 539](#).



Caution: Do not exceed more than 10,000 resources in a group.

Where: **Navigator > Resources > Notifications > Destinations > Shared > All Destinations**

To create or edit destination groups:



Note: As a user, you can create new groups under **All Destinations**, but not new subgroups under existing system-defined groups.

1. If you are creating a group, right-click **All Destinations** and select **New Group**.
If you are editing a group, right-click the group and select **Edit Group**.
2. Set the following attributes:

Field	Description
Name	Enter a name for the group. Caution: If you are renaming an existing group, check for the following: • If this group is assigned to a user, edit that user's Notification Groups attribute by selecting the renamed destination. • If a rule action to send notification is set to send to this group, edit that rule action by selecting the renamed destination.
Reference Pages	Enter text as required for the group or for members. Refer to "Reference Pages" on page 1027 for information.

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

3. Press **Enter**.

To rename destination groups:



Caution: If you are renaming an existing group, check for the following:

- If this group is assigned to a user, edit that user's Notification Groups attribute by selecting the renamed destination.
- If a rule action to send notification is set to send to this group, edit that rule action by selecting the renamed destination.

1. In the **Notifications** resource tree, right-click a group and choose **Rename**.
2. In the text field, rename the group.
3. Press **Enter**.

To delete destination groups:



Caution: If you are deleting a destination group, check for the following:

- If this group is assigned to a user, edit that user's Notification Groups attribute by selecting the another destination.
- If a rule action to send notification is set to send to this group, edit that rule action by selecting another destination.

1. In the **Notifications** resource tree, right-click a group and choose **Delete Group**.
2. Click **Yes**.

Managing Notification Destinations

Destinations are mapped to user groups, therefore, make sure the user group exists (["Managing User Groups" on page 99](#)) for the destination you are creating.

To create or edit destinations:

1. In the Notification resource tree in the Navigator panel, right-click an escalation level (such as **Level 1**).
2. If you are creating a destination, choose **Add New Destination**.

If you are editing a destination, right-click a notification destination and choose **Edit Destination**.

For more information, see ["Changing Notification and Acknowledgment Settings" on the next page](#).

3. In the Notification Editor, enter a label for the notification in the **Name** field.
4. Optionally set a **Start Time** and **End Time** during the day within which the notification is to be active. The default is all day (12:00:00 AM to 11:59:59 PM).
5. Select a **Destination Type** and the related parameters for that type, as follows:
 - For **Console**, additionally select a user or a user group. This displays the notification on the users' ArcSight when they log in.
 - For **Email Address**, enter a valid email address and additionally select a user or user group.



Note: Always set the ArcSight **User/Group** identity for all destination types.

6. For **User/Group**, select the group from the resource selector popup.

7. (Conditional) If you are adding more than one destination to an escalation level, use a unique user group for each destination. Using a group to which another destination under that escalation level belongs might result in incorrect notification numbers.
8. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
9. Click **OK**.

To move or copy destinations:

1. In the Notification resources tree, find a destination and drag it to a different escalation level. You can drag across groups if needed.
2. Right-click the destination and choose **Move** to move it, **Copy** to make a separate copy, or **Link** to create a copy of the destination that is linked to the original destination.

If you choose **Copy**, you create a separate copy of the destination that is not affected when the original destination is edited. If you choose **Link**, you create a copy of the destination that is linked to the original destination. Therefore, if you edit a linked destination, whether the original or the copy, all links are edited as well. When deleting linked destinations, you can either delete the selected destination or all linked destination copies.

To delete destinations:

1. In the Notification resource tree, right-click a notification destination and choose **Delete Destination**.
2. In the dialog box, click **Yes**.

Changing Notification and Acknowledgment Settings

Administrators can configure notifications, acknowledgments, and wait-time settings. The escalation time window or wait-time depends on the event's severity.



Note: If notifications and acknowledgments were disabled during Manager setup, mail server settings made through the ArcSight Console do not take effect until you re-run the Manager configuration wizard to enable notifications and acknowledgments on the Manager side.

To re-enable notifications settings, follow instructions in the [ESM Administrator's Guide](#) to:

- Stop the ArcSight Console and Manager.
- Re-run the Manager configuration wizard and enter settings for your SNMP server.

To change e-mail settings:

1. In the Notification resource tree, right-click a group and choose **Settings**, then **Edit E-mail Settings**.
2. In the Notification Editor, type in the following text fields:

Notification Fields	Definition
From Address	The e-mail address from where the notification messages are sent. It is important that the "from address" specified is one that is not rejected by the SMTP server, since some SMTP servers reject unknown e-mail addresses. For notifications sent by cell phone, any cell phone must be e-mail enabled.
Outgoing Mail Server	The host name of the local outgoing mail server. This is the SMTP server ArcSight uses to send e-mail. The Outgoing Mail Server must be accessible from the ArcSight Manager for e-mail notifications to be sent. SMTP is used to send e-mail. An SMTP server must be configured either at install time or set here.
Incoming Mail Server	The local incoming mail server host name.
Incoming Mail Protocol	Select either IMAP or POP3 mail protocols.
E-mail Account	The e-mail account name. For notifications sent by e-mail, you need to add an address to the e-mail Address field.
Account Password	Enter the password for the account.
Confirm Password	Re-enter the same password to confirm.



Note: POP3 and IMAP can be used to check for e-mail acknowledgments. You can specify these options at install time, or set them here. For acknowledgements, the relevant fields are "incoming mail server," which is the POP/IMAP server to specify to check e-mail, "incoming mail protocol," which is either POP3 or IMAP, "account" and "password," which are the login name and password to access the mailbox from the incoming mail server. Note that replying to mails from the notification "from address" should reach the mailbox accessible to the "account" login.

3. Type the **E-mail Account** password in the Password text field and confirm it in the Confirm **Password** text field.
4. Click **OK**.

To change wait time settings:

The default wait-time values for Very-High severity and High severity are set at 5 minutes, Medium is set for 30 minutes, and Low is set for 2 hours.

1. In the Notification resource tree, right-click a group and choose **Settings**, then **Edit Escalation Wait Time**.
2. In the Notification Editor, type in the wait time for the hour (**Hr**) and minute (**Min**) text fields for **Very-High**, **High**, **Medium**, or **Low** severity.
3. Click **OK**.

Testing Notification Groups and Destinations

This topic describes how to test notification groups and destinations.

To test group notifications:

In the Notification resource tree, right-click a populated notification group and choose **Test Group Notification**.

A test notification message is sent to the notification destination. Test notifications are not sent to group notification destinations if the End Time has expired. For example, if you test group notification at 6:00:00 PM and the End Time states 5:00:00 PM, a notification message is not sent to the group.

To test destination notifications:

In the Notification resource tree, right-click a notification destination and choose **Test Destination Notification**.

A test notification message is sent to the notification device. Test notifications are not sent to notification destinations if the End Time has expired. For example, if you test a notification destination at 6:00:00 PM and the End Time states 5:00:00 PM, a notification message is not sent to the device.

Managing Escalation Levels

This topic describes how to handle the tasks required for managing escalation levels.

To add an escalation level:

In the Notifications resource tree, right-click a notification group and choose **Add Escalation Level**.

New escalation levels are added in sequential order. If you want to add a level between two existing levels, add another level then move destinations accordingly. For example, if you have

Level 1 and **Level 2** and you want to add a level between them, add another level, **Level 3**. Then, move all destinations from **Level 2** to the new **Level 3**.

To delete an escalation level:

1. In the Notifications resource tree, select the last escalation level in a notification group.



Note: All destinations within this escalation level are also be deleted. If you want to save the destinations, make sure you move them to another level **before** deleting.

2. Right-click the escalation level and choose **Delete Escalation Level**.

Chapter 8: Monitoring Events

This topic describes how to monitor events coming from SmartConnectors using tools that are displayed in the Viewer panel. You can monitor events through a rich set of views, including active channel and grids, dashboard graphics and tables, and active lists, as described in the following topics:

Monitoring Active Channels

Active channels provide a streaming view of events coming into your system that can be viewed numerous ways using numerous types of filters and field sets.

Creating or Editing an Active Channel

This topic shows how to create active channels manually, from triggered rules, and from filters.



Tip: Press **Enter** to register edits made in editors and channel columns.

To ensure that ESM registers a change you make to a field in editor and channel columns, press **Enter** before clicking **Apply** or **OK**.

Where: Navigator > Resources > Active Channels

To create or edit an active channel:

1. Locate an active channel group.
2. If you are creating an active channel, select **New Active Channel**.
If you are editing an active channel, expand the group, right-click an active channel, and choose **Edit Active Channel**.
3. Set these options:

Active Channel Attributes

Attribute	Usage
Start Time	The relative or absolute time reference that begins the period to track events in the channel. Edit the time expression, choose a common expression from the drop-down menu, or click the Selector button to choose an absolute date and time value. See "Timestamp Variables" on page 1060 for more options. Notes: • Time intervals: Based on Start Time and End Time settings, the active channel rounds up time as follows: If the interval from start to end time in minutes is less than 10 minutes, the time interval will be at seconds precision. If the interval from start to end time in minutes is greater than 1440 minutes, then the time interval will be at hours precision. Otherwise, the time interval will be at minutes precision. • Change in Daylight Savings Time: If a channel is open when Daylight Savings Time starts or ends, it does not show the correct start time until you restart it. You can change the default start time for new channels by editing the console.properties file in the <ArcSight_Console_HOME>/ current/config directory. For example, add the this line... <code>console.channel.newChannel.defaultSubtractTime="\$Now - 2h"</code> ... to change the start time to two hours ago. For a list of possible time values see the Start Time: field pull-down menu. Refer to "Managing and Changing Properties File Settings" in the ESM Administrator's Guide.

Active Channel Attributes, continued

Attribute	Usage
End Time	The relative or absolute time that ends the period to actively track the events in the channel. Edit the time expression, choose a common expression from the drop-down menu, or click the Selector button to choose an absolute date/time value. See "Timestamp Variables" on page 1060 for more options. Notes: • Time intervals: Based on Start Time and End Time settings, the active channel rounds up time as follows: If the interval from start to end time in minutes is less than 10 minutes, the time interval will be at seconds precision. If the interval from start to end time in minutes is greater than 1440 minutes, then the time interval will be at hours precision. Otherwise, the time interval will be at minutes precision. • Change in Daylight Savings Time: If a channel is open when Daylight Savings Time starts or ends, it does not show the correct start time until you restart it. If a channel is open when Daylight Savings Time starts/ends, the live channel does not show the correct start time until you restart it. • If setting the End Time results in the message "Invalid end date for sliding channel," the channel is set to <code>Continuously evaluate</code> instead of <code>Evaluate once at attach time</code>. Either re-set the End Time or change the Time Parameters option for the channel to <code>Continuously evaluate</code>. • Avoid creating active channels that query more than once day. For active channels that query more than once day, use <code>Evaluate time parameters once at attach time</code> instead of <code>Continuously evaluate</code>. Better yet, use trends for these types of active channels. See also "Best Practices to Optimize Channel Performance" on page 246.
Use as Timestamp	Choose the event-timing phase that best supports your analysis. End Time represents the time the event ended, as reported by the device. Manager Receipt Time is the event's recorded arrival time at the ArcSight Manager.

Active Channel Attributes, continued

Attribute	Usage
Evaluation of time parameters	Choose whether the channel will Continuously evaluate to show events that are qualified by Start and End times which are re-evaluated constantly while the channel is running, or Evaluate once at attach time to show only the events that qualify when the channel is first run. A channel set to Continuously evaluate is also known as a <i>sliding channel</i> , and typically has its End Time option set to \$Now .
Filter	If creating a new channel, select an existing filter for the events processed through the channel. If you prefer, click Define to create a new filter to be used by this channel. Follow the instructions in If editing a channel, go to the Filter tab to make your edits.
Fields	Choose an existing event field set for the events processed through the channel. The default field set is for users who view a channel for the first time. If no default is specified, the ArcSight system default is used. When a user closes a channel, ArcSight saves the field set (and all other Console settings) to the user's .ast file. After a user has opened a channel once, the Console does not use the default field set for that user again. Changing the default only affects other users who have never opened the channel before.

- Click the **Examples** button to see how to specify commonly used channel values.

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

- Click the **Filter** tab to edit the channel's filter condition as described in ["Creating or Editing a Filter" on page 306](#).

To view the full conditions for the **MatchesFilter** operator, click the **Summary** tab and then click the **Expand Filter** button to display the filter conditions for debugging.

Note that in this case, the display of the **MatchesFilter** full logic does not display the sub-filter of the matched filter. Full logic is displayed only for the first level of matched filter conditions.

- Click the **Sort Fields** tab to explicitly set which fields to sort the channel on in grid views, the sort order for those fields, and whether sorting for each field is ascending (A to Z) or descending (Z to A).
- Click the **Local Variables** tab to use ArcSight local variables with the channel's filters.



Tip: You can create local variables, which are only available to the resource you are creating (in this case, an active channel), or use global variables. For information on creating global variables, see ["Creating or Editing a Filter" on page 306](#) and ["Global Variables" on page 609](#).

8. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
9. Click **OK** to save the channel and to open and run it in the Viewer panel.

To view results of triggered rules in channels:

See ["Verifying Rules with Events" on page 553](#).

To create active channels from filters:

- In the Filters resource tree, right-click a filter and select **Create Channel with Filter**.
- Do the same for:
 - Connectors
 - Assets, including vulnerabilities, zones, and categories
 - Stages
 - Cases with events. For such cases, right-click and select **Case Details Channel**. See ["Viewing a Case's Events in a Channel" on page 643](#).

Viewing Active Channels

Viewing and using active channels include creating them, filtering them, customizing contents, changing presentation formats or layouts, and deleting them.

When viewing an active channel, keep the following in mind:

- If a channel is open when Daylight Savings Time goes into or out of effect, the live channel will not reflect the correct start and end times until the channel is closed and re-opened.
- If an active channel uses a filter that applies conditions to a list data type field, the active channel will include multiple rows for the same event or resource (for example, actors) channels. This behavior is expected.

Purpose:

- View events
- View resources such as assets, vulnerabilities, cases, and so on

Where: **Navigator > Resources > Active Channels**

To view an active channel:

1. Right-click a channel.
2. Select **Show Active Channel**.

The selected channel is displayed in the Viewer.

To edit the time window for an active channel:

1. Click the edit time window button that looks like a pencil (next to the active channel Start Time).
2. Modify the **Start Time** and **End Time**.
3. Click **OK**.

To view resources in active channels:

Right-click a resource or group, and choose **Show** <ResourceName>. The resources are displayed in an active channel view.

To view results of triggered rules:

See ["Verifying Rules with Events" on page 553](#).

Monitoring Events in the Active Channel

Click an active channel's tab at the top of the Viewer panel and select the **Grid** view of that channel using the tab at the bottom. When new events occur, they are displayed at the top of an active channel as a new row. Events can appear in ArcSight Severity or filter colors. You can set the color-code for events by using the steps described in ["Setting Grid Options for the Viewer Panel" on page 52](#).

Full Search and Event Search on ArcSight Command Center

You can launch a search from the Console's event channel to the ArcSight Command Center. The following types are supported:

- **Event Full Search** is searching for a matching value (for example, abc) stored in multiple fields.
- **Event Search** is searching for a matching value (for example, abc) stored in one specified field (for example, only in the Name field).

To launch a search:

1. Open an active channel.
2. Select a cell with the value you want to search. For example, the field (column) is Target Address and its value is 198.51.100.0.

3. Right-click the cell that has the value. From the context menu:

- To search for the value 198.51.100.0 on multiple fields (not just the Target Address field), select **Event Full Search**.
- To search for the value 198.51.100.0 on a specific field on the channel, for example, Target Address, select that cell, select **Event Search**.

The ArcSight Command Center is launched on the preferred browser.

4. Enter your name and password to log in.

The Command Center displays the Event Search page with the results.

- If you selected Full Event Search, the search results will display instances of 198.51.100.0 in the source address and destination address fields.
- If you selected Event Search, the search results will display instances of 198.51.100.0 in the Destination Address field only, because you came from the Target Address field on the Console's active channel.

Refer to the [ArcSight Command Center Users Guide](#) for more information about the search feature.

Supported fields for the search are:

Supported fields		
AGENT_ADDRESS	AGENT_HOST_NAME	AGENT_NT_DOMAIN
AGENT_SEVERITY	AGENT_TYPE	AGENT_ZONE
AGENT_ZONE_NAME	AGENT_ZONE_RESOURCE	AGENT_ZONE_URI
APPLICATION_PROTOCOL	BYTES_IN	BYTES_OUT
CATEGORY_BEHAVIOR	CATEGORY_CUSTOM_FORMAT_FIELD	CATEGORY_DEVICE_GROUP
CATEGORY_OBJECT	CATEGORY_OUTCOME	CATEGORY_SIGNIFICANCE
CATEGORY_TECHNIQUE	CATEGORY_TUPLE_DESCRIPTION	CUSTOMER
CUSTOMER_NAME	DESTINATION_ADDRESS	DESTINATION_DNS_DOMAIN
DESTINATION_HOST_NAME	DESTINATION_MAC_ADDRESS	DESTINATION_NT_DOMAIN
DESTINATION_PORT	DESTINATION_PROCESS_NAME	DESTINATION_SERVICE_NAME
DESTINATION_TRANSLATED_ADDRESS	DESTINATION_USER_ID	DESTINATION_USER_NAME
DESTINATION_USER_PRIVILEGES	DESTINATION_ZONE	DESTINATION_ZONE_NAME
DESTINATION_ZONE_RESOURCE	DESTINATION_ZONE_URI	DEVICE_ACTION
DEVICE_ADDRESS	DEVICE_CUSTOM_NUMBER1	DEVICE_CUSTOM_NUMBER2
DEVICE_CUSTOM_NUMBER3	DEVICE_CUSTOM_STRING1	DEVICE_CUSTOM_STRING2

Supported fields		
DEVICE_EVENT_CLASS_ID	DEVICE_PRODUCT	DEVICE_VENDOR
DEVICE_VERSION	DEVICE_DOMAIN	EVENT_ID
EXTERNAL_ID	FILE_NAME	FILE_PATH
FLEX_NUMBER1	FLEX_NUMBER2	FLEX_STRING1
FLEX_STRING2	GENERATOR	LOCALITY
MESSAGE	NAME	OLD_FILE_HASH
OLD_FILE_ID	OLD_FILE_NAME	OLD_FILE_PATH
OLD_FILE_PERMISSION	OLD_FILE_SIZE	OLD_FILE_TYPE
REASON	RELEVANCE	REQUEST_CLIENT_APPLICATION
REQUEST_CONTEXT	REQUEST_METHOD	REQUEST_URL
REQUEST_URL_FILE_NAME	REQUEST_URL_QUERY	SESSION_ID
SEVERITY	SOURCE_ADDRESS	SOURCE_DNS_DOMAIN
SOURCE_HOST_NAME	SOURCE_MAC_ADDRESS	SOURCE_NT_DOMAIN
SOURCE_PORT	SOURCE_PROCESS_NAME	SOURCE_SERVICE_NAME
SOURCE_TRANSLATED_ADDRESS	SOURCE_USER_ID	SOURCE_USER_NAME
SOURCE_USER_PRIVILEGES	SOURCE_ZONE	SOURCE_ZONE_NAME
SOURCE_ZONE_RESOURCE	SOURCE_ZONE_URI	TRANSPORT_PROTOCOL
SOURCE_ZONE_URI	VULNERABILITY_EXTERNAL_ID	VULNERABILITY_URI

Using Views

Views can vary in scope and scale, from broad to narrow, and from graphic to detailed, depending on how your enterprise is organized and monitored.

To select a view:

In the Viewer panel, click a tab at the **top** to choose an active channel by name. On a channel, you can select various instances of that channel (such as a grid view and bar chart of the same data) by clicking its tile or its tab at the **bottom** of the panel.

Alternately, to advance quickly through each of the tabs in the Viewer panel, press **Ctrl+Shift+N** or **Ctrl+Shift+A** to jump forward or backward, respectively. These keystrokes apply to any type of view in the Viewer panel.

To change view layouts:

You change individual view layouts with the **Layout Selector** menu available from the blue icon at the lower-right corner of the Viewer panel. Click this icon to choose:

Layout Option	Result
Tab	Fill the active channel display with the current view and make other open views selectable by tabs at the lower border.
Tile Best Fit	Display all views in the active channel as variously shaped tiles, giving each a proportional amount of space.
Tile Horizontally	Display all views in the active channel horizontally, giving each a proportional amount of space.
Tile Vertically	Display all views in the active channel vertically, giving each a proportional amount of space.

To float a view:

In the active channel's name tab, right-click and choose **Float**.

To close one or all views:

In the active channel's name tab, right-click and choose **Close** or **Close All**.

To close an individual view **Shift+click** its name tab. You can also right-click a view name tab and choose **Close** from the popup menu.

To close all views except the current one:

In the active channel's name tab, right-click and choose **Close All But Current**.

To add another type of presentation (view) for the data in an active channel:

1. Click the **View Type** icon in the lower-right corner of the Viewer panel.
2. Select from among the grids and the various types of chart or graphic views.

To change view layouts of channels within a container:

This technique applies to individual channels within a view container, such as data monitors within a dashboard.

1. Click the **Layout** icon
2. Select to show or arrange the views by **Tab**, or **Tile Best Fit**, **Tile Horizontally**, or **Tile Vertically**.

Investigating Views

This topic explains how to use the Console's **Analyze in Channel** command to refine and explore channels contextually, using attributes of the events already being displayed in grid views.

The **Analyze in Channel** command uses these attributes, and the values found in their events, to automatically formulate simple filters or conditions.

When you create or refine a filter through **Analyze in Channel**, the Viewer panel automatically opens a new view of the channel with the filter applied. You explore the filter's effect in this view. You can keep the view by saving the channel under a new name, or discarding it by right-clicking in the grid and choosing **Close**.

When you use **Analyze in Channel** to add a condition to a resource editor such as Rules or Filters, the condition appears in the editor panel where you can modify it or click **Apply** to put it into effect.

The new or modified views you generate with the **Analyze in Channel** command can be grids, or you can choose to display them in applicable chart formats using the **Viewer Selector** icon in the lower-right corner of the Viewer panel.

To learn more about the event attributes these options use, see ["Data Fields" on page 910](#).

Viewing an Exploited Vulnerability

The **Analyze In Channel** options include the ability to look for potentially exploitable vulnerabilities associated with an event.

1. Select an event in a grid view.
2. Right-click the event and choose **Analyze in Channel > Show Exploited Vulnerability**. Available information appears in the Vulnerabilities tab of the relevant Asset Editor.

Viewing a Targeted Asset

You can also find out more about an asset targeted by an event.

1. Select an event in a grid view.
2. Right-click the event and choose **Analyze in Channel > Show Targeted Asset**. Available information appears in the Asset Editor.

Filtering an Active Channel

You can filter active channels through the Filter tab of the Active Channel Editor.

1. Right-click the filter name in the header.
2. Select **Edit Filter** to open the editor.
3. Create a filter as described in ["Creating or Editing a Filter" on page 306](#).



Tip: Understanding how to use the Common Conditions Editor (CCE) is integral to creating and editing filters. See ["Common Conditions Editor \(CCE\)" on page 887](#) for more information.

See also ["Filtering Active Channels with Inline Filters" below](#).

Filtering Active Channels with Inline Filters

Active channels have a means for creating simple filters based on using a value found in one column, or creating AND conditions between values found in two or more columns.

These filters are called *inline filters*, a very rapid way to constrain detailed views. While inline filters are in use, they affect all views generated for the channel.


You can create, change, save, hide, and remove inline filters from the active channel. Also, you can create and manage multiple inline filters from this view.



Note: Custom columns are not available as arguments for inline filtering.


To create an inline filter:

1. Click the **Inline Filter** link in the event header or click the **Edit Inline Filter**  button at the top right of the active channel to display the inline filtering fields.



The screenshot shows the 'Active Channel' configuration panel. At the top, it displays 'Total Events: 4,220' and a severity distribution: Very High: 0, High: 0, Medium: 0, Low: 4,220, Very Low: 0. Below this is a 'Radars' section with a bar chart. The main section is the 'Inline Filter' configuration, which is circled in red. It includes a table with columns: Manager Receipt Time, Name, File Name, Target User Name, Simple Attacker Target, and Category. Below the table are buttons for 'Highlight', 'Apply', 'Cancel', 'Clear', and 'Help'.


2. Enter a value by which you want to filter for one or more fields relating to a column in the grid:
 - a. Click inside a field to display an equals operator and a selector button to display a list of possible values for the field.



These possible values are based on the displayed events.

- b. If desired, click the ellipsis button (...) to bring up the Condition Editor dialog in which to create the filter for the selected field.
3. Click **Apply** to apply the filter immediately to the view. The inline filter is displayed in the header under the standard filter.

The channel header displays the inline filter name under the standard filter used by the channel.



To change an inline filter:

1. Click the **Edit Inline Filter** button.
2. Optionally click **Clear** if you want to remove the old fields.
3. Choose new values and click **Apply**.

To remove an inline filter:

Right-click over the Inline Filter name in the header for the selected event and choose **Remove Inline Filter**.

To save an inline filter:

1. Right-click over the Inline Filter name in the header for the selected event.
2. Choose **Save Inline Filter**.

This opens a Filters Selector dialog that shows the Filters tree.

3. Navigate to the Filters subfolder where you want to save the current filter, and click **OK**.

To highlight the filtered events with color:

1. Click the **Highlight** check box (**on** is check marked).
2. In the dropdown, select a color from the palette.

To create and manage multiple inline filters:


- Click the + button next to the Highlight options under the inline filters to add filter definition rows.
- Click the - button to remove filter rows.

Multiple inline filters provide a means of creating a filter with complex conditions, inline in an active channel. For example, in the Name column for an event, you could specify that the event name contains `ActiveList` on the first filter row and that the name does not contain `Successful`. You could extend this filter by specifying what you are looking for in some of the other fields or even add more qualifiers on the Name field. All fields can be narrowed down in this way, using multiple filter definition rows.

Applying a Field Set to an Active Channel

To apply a field set to an active channel:

1. Right-click over any field header and choose **Field Sets > Select Field Set** to open the Field Sets Selector dialog.
2. In the Field Sets Selector dialog, select a field set and click **OK**. Note that domain field sets apply to Oracle-based ESM.



The active channel is displayed with the selected field set.



Note: About ArcSight System Sortable Field Sets


The Sortable Field Sets under ArcSight System are not available for selecting in active channels. The ArcSight System sortable field sets are a special set marked for internal use to provide the sortable functionality and maintain consistency between the Console user interface, field sets, and database indexes.

For more information about sorting, see ["Sorting Events in the Active Channel" on the next page](#).

See ["Variable Availability and Contexts" on page 1087](#) for information about using variables in active channels.

Using an Active Channel Header

Each active channel has a header section with several features you can use to understand and manipulate what the channel displays.



Active Channel Header Features

Feature	Usage
Name and Total	The top line of the header shows the channel's name and how many events it contains. You can also use the Minimize or Maximize button at the right end to close or open the header. Note: The event count function on Active Channels only reports live events, not replay events. If you prefer to see a count of all events coming through during a particular period, you should create a query viewer or report. If you want a count of only replay events, the event count in a replay channel will provide an accurate count of all replay events within a specific time window.
Priority Indicators	On the right border of the header is a column of event-priority statistic indicators. The numbers beside the Priority categories are the number of events in those categories. Click these indicators to filter the channel to show only the selected priority.
Time Span	The Start Time and End Time show the chronological range of the channel.

Active Channel Header Features, continued

Feature	Usage
Filter status	This describes the filter that limits what the channel shows. Click a filter status name, such as <No Filter> , to open the Active Channel Editor and its Filters tab, where you can add, edit, or delete contents as described in "Creating or Editing a Filter" on page 306 . You can also right-click the current filter status and choose to edit, save, or remove it.
Radar display button	Close or open the display with the Minimize or Maximize button at the right end of the Filter line.
Radar display operation	Click, Shift+click, Ctrl+click , or drag to select bars in the display. You can also drag a selection's borders left or right. The grid then shows just the events the selection represents. The display shows This channel is active but temporarily empty at any time, no matter how briefly, if there are no qualifying events. This also might show when a channel first opens.

Sorting Events in the Active Channel

The names of sortable fields in column headers are indicated with a double-arrow icon . If a field is already sorted, an up  or down  arrow indicates the direction of the sort. If the column contains words or alphabetic characters, it sorts alphabetically from A to Z (or vice versa).



Note: You might experience performance problems when sorting columns in an active channel. Some columns are resource-intensive to sort, such as string fields containing 1000 characters. Consider using query viewers instead, where you have the option to group and order fields.

See ["Best Practices to Optimize Channel Performance" on page 246](#) for additional information.

To sort events in an active channel:

- To sort the list by a column, right-click over the column and select **Sort Column**.
- To reverse the sort order, select **Sort Column** again on an already-sorted column. This makes the column the primary sort column.
- To remove a sort, right-click over a sorted column and select **Remove Sort**.

You can also perform an advanced sort on one or more columns in the active channel. When selecting a secondary sort column, select the secondary column first, then the primary column. For example, to sort by Event Name then by Detect Time, sort **Detect Time** first, then **Event Name**.

After you sort a column it automatically pauses the current channel, stopping events from appearing in the active channel. Click the **Play** button in the Replay Controls to restart the channel and resume receiving events in the active channel.



Note: When you sort on time and on priority, you might observe cases where events with the same apparent time are not in priority order. Because events are timestamped to milliseconds, they may in fact be in time order although the milliseconds are not showing. In this case, you can show milliseconds to validate time order. Choose **Edit > Preferences**, then in the Date and Time panel change the **Date & Time** Format to also show milliseconds by adding SS to the seconds parameter, for example, d MMM yyyy HH:mm:ss:SS z.

For additional information, see ["Applying a Field Set to an Active Channel" on page 236](#).

About primary and secondary sort columns:

When you sort a column, it becomes the primary sort column and the number 1 appears next to the sort arrow. The previous column by which the report was sorted becomes the secondary sort column and the number 2 appears next to the sort arrow.

This numbering applies to every column on which you sort: the newest sort is always 1 and the others change accordingly.

To sorting by primary and secondary time columns:

When your primary and secondary sort columns are both time columns (such as Create Time and Modification Time), milliseconds become a factor in sort order.

Milliseconds are not displayed. This can create a situation where a number of items with the same primary time appear to show the secondary sort as in the wrong order. In reality the primary times are off by milliseconds, so they are *not* the same, and these milliseconds affect sort order before the secondary time is taken into account.

Adding, Replacing, or Removing a Column


A quick way to add, replace or remove columns in a active channel (for example, active channel or list) is to right-click on the appropriate column header and select one of the following options:

- **Columns > Add/Remove Column > <Select a field from the menu>**
- **Columns > Replace This Column > <Select a field from the menu>**
- **Columns > Remove This Column**

These are context-dependent commands that apply to the column on which you launch the right-click menu. (To add a column, right-click on the header of the column you want to add the new column next to. Columns are added to the right of that header. To replace or remove a column, right-click on the header of the column you want to replace or remove.)

Alternatively, you can use the Customize Columns dialog to define the columns shown in the viewer as described here:

1. Right-click the column header and select **Columns > Customize Columns** to bring up the associated dialog. (Fields shown in *italics* are *derived* fields.)



The dialog is an example of what is displayed, based on the columns on the channel.



Tip: Looking for information about custom columns? If you want to add a custom column, you need to create (define) it first. Once it's created, it appears in the Available Fields list under Custom Column, and you can include it in active channels as with any other field. For information on creating custom columns, see ["Customizing Columns" on page 250](#).

- **To add a column:** Select data fields (column titles) to add from the Available Fields list on the left. Check marks indicate selected columns. The selected columns show up in the list on the right as you select them. Alternatively, when you deselect or uncheck a data field on the left, the column is removed from the right-hand list.
- **To remove a column:** Select a field from the right-hand list and click the Delete button . Also, deselecting a data field from the Available Fields list on the left removes it from the right-hand list. Removing a column from an active channel does not delete the column information from the database.



Tip: You also can remove a column directly from the active channel without opening the Add/Remove Columns dialog. To do this, right-click a column header and select **Remove Column**.

- **To re-order the columns:** Select a data field (column title) in the right-hand list and click the Up  and Down  buttons to move it. The top-to-bottom order shown in the Show columns in this order list on the right translates to a left-to-right order when applied in the active channel. A column title at the top of this list will

show as the first column in the channel on the far left in the grid display. A column title at the bottom of this list will show as the last column on the far right of the grid.

2. Click **OK** to save changes you made on the Add/Remove Columns dialog. The active channel reflects added, replaced, removed, or re-sorted fields.

Sizing, Showing, or Hiding Column Elements

To change the column size:

Right-click a column header and select **Size Column To Fit**.

To showing or hide column text and icons:

Right-click a column header and select one of the following options:

Option	Display Result
Text and Icon	Display the column heading and its icon.
Text Only	Display only the column heading.
Icon Only	Display only the icon.

Using Active Channel Menu Commands

Right-click an event or event field in the active channel to open a context menu. The commands available are those that apply to the current combination of event type, view, filter, and so forth.

Active Channel Menu Commands

Command	Description
Show Event Details	Use the Event Inspector to examine all the attribute details associated with the event.
Correlation Options	• Simple chain: Show this event's base and correlated event tree in the Event Inspector. • Detailed chain: Show this event's base and correlated events in detail in a new active channel. • Show triggering resource: Show the rule that triggered this event in the Rule Editor. • Clear rule actions: Clears the list (if one is showing) of rule actions pending on the database.

Active Channel Menu Commands, continued

Command	Description
Analyze in Channel	Create a temporary filter as required based on the field's highlighted event. The Analyze in Channel command uses the event's attribute type (its column heading), and the particular event's field value (for example, an exact IP address), to formulate simple filters based on these two factors. The filter's operators can include Create Filter [X = Y] and Add Condition [X = Y] to Editor . The Analyze in Channel submenu also offers the Show Exploited Vulnerability and Show Targeted Asset commands to open detailed views of assets or vulnerabilities, if present in the selected event.
Debug Filter	Evaluate if the selected event passes the filter resource selected from the filter resource popup.
Debug Event Priority	Display information on how event priorities are determined for the selected event. The window lists which conditions match the event. Items under each category: Severity, Relevance, Model Confidence, and Asset Criticality, and the total scores. For each category, certain factors contribute their individual scores. The scores are added to calculate the total. However, if the sum exceeds the upper limit of 10, 10 is displayed for the category's total score. Lower limit is 0. Debug Event Priority is applicable to Threat Level Monitoring, described in "Threat Evaluation" on page 1056 and also "Priority Calculations and Ratings" on page 1019.
Active List	Add the selected event to, or remove it from, an active list. See "Adding Events from a Channel to an Active List" on page 497 .
Annotate Event	Open this event in the Annotate Events dialog box, where you can click the Stage field to set a collaboration workflow sequence for this event. When you select a stage you automatically place the event in the corresponding group in the Stages resource tree in the Navigator panel, where you and other analysts can collaborate on its investigation and resolution.
Mark as Reviewed	Set the event's annotation flag to IsReviewed. See "Event Annotation Group" on page 939 , specially the Flags label, for a list of event annotation fields.
Event Graph	Graph any logical relationships (that is, source/target IP address connections) that exist among the currently selected events.
Rule Chain Graph	Graph the rule chains behind the currently selected triggered events.
Geographic View	Geographically map the source and destination IP addresses of the selected events.
Integration Commands	Link to other ArcSight applications and tools. For more information, see "Integration Commands" on page 664 .
Tools	Displays the Tools command menus (also available from the menu, Tools > Local Commands . See "Using the Network Tools" on page 73 and "Using the Tools Menu" on page 84 .
Export	Export the selected events to an external event-tracking system, such as comma-separated-value (CSV) data in a report or for a spreadsheet, or save it as an HTML or a JPEG file.
Add to Case	Add the selected events to a new case for tracking.
Payload	Keep or discard the payload associated with a selected event. Disabled if the selected event has no associated payload.

Active Channel Menu Commands, continued

Command	Description
Report	Includes two options: • Event Context Report: Output a report concerning rules and events within a specified time window. See "Running an Event Context Report" on page 452 for details. • Channel Report: Output a report of all events in the channel. See "Running a Channel Report" on page 452 for details.
Close	Close the current individual view within the selected view type.
Knowledge Base	If defined, show the Knowledge Base pages associated with the selected events, or associate new pages.
Reference Pages	If defined, displays the reference pages for this event.
Vendor Page	If available, show vendor Web page of the event's sensing device.
Help	Open the online Help to this topic.

Exporting Events to a File

You can export a set of data fields into a comma separated values (CSV) file.

The procedures in this topic instruct you to:

- Export data fields into a CSV file
- Limit exported fields to only those fields that are visible on the channel

To export to a CSV file:

1. In the channel, select one or more events.
2. Right-click and select **Export > Events in Channel**.
3. On the Export Events file browser, navigate to the location where you want to save the

CSV file, then enter or select options for these fields:

File Name	Enter a file name for the CSV file. Note: The file name extension is not required; the csv extension is added automatically when the file is created.
Files of Type	Select Comma separated values (*.csv).
Export Data Options	For Rows, you have two options: • If you select All in channel, all events in the channel will be exported to the CSV file. • If you select Selected rows only, only those rows highlighted for the right-click operation will be exported to the CSV file. • The default for Columns is the Export field set. Keep the default or select other field sets from a list of All Field Sets. • For Destination, select Local CSV File.

4. Click **OK**.

To limit the export to fields visible in the channel:

Purpose: The default “Export” field includes a large number of columns. Unless you have a pressing need to export all these fields for channel events, you might want to modify the export. Exporting a large field set for a large event set could be time- and resource-consuming.

- If the channel is unmodified from its default (you have not added or removed fields), you can select the channel’s default field set on the file export option. To find the default field set name, edit the channel and look at “Default Field Set” name or right-click any column header in the channel and choose **Field Set > Selected Field Set**. The default field set will be selected. (For example, for /All Active Channels/ArcSight System/Core/Live active channel, the default field set is Standard-MgrRcpt. Selecting this field set on the export will give you that set of columns in the CSV file.)
- If you have modified the channel from its default (added or removed fields), you can save it as a custom field set and then choose your custom field set on the export dialog. To save a custom field set, right-click anywhere on the column headers in the active channel and choose **Field Sets > Save As**. On the Field Sets Selector, navigate to the group you want, name the new field set and click **OK**. Now it will be available to choose from on the export dialog.

The Export field set itself is also customizable. If you are sure you always want exported events to include a limited set of fields, you can edit the Export field set. See ["Creating a Field Set" on page 601](#) and ["Editing a Field Set" on page 607](#).

Defining Grid Fields Options

In the New Active Channel dialog box you can choose from the **Select a Field Set** menu, or you can click the **Define** button to open the Define Grid Fields dialog box. See ["Creating a Field Set" on page 601](#) for more information. To change these choices after creating a channel, use the steps described in ["Customizing Columns" on page 250](#).

Grid Field Options

Option	Usage
Fields	A name for the set
Available Fields	Select the event fields (also called data fields or attributes) that you want the channel to process. As you make selections, they appear in the Fields to Show list at the right. Remember that not all fields are readily sortable.
Fields to Show	This list shows the selections you have made in the Available Fields list. The order you give to the fields in this list becomes their default presentation order in grid views. Once populated, you can select one or more fields (Shift+click and Ctrl+click apply) to rearrange with the Move Up , Move Down , and Remove buttons.
Move Up, Move Down, Remove buttons	These buttons move or remove the fields you select in the Fields to Show list. The order you set becomes the presentation order in grid views.
Sort First By	After selecting and ordering fields, you establish their sorting order (also called their <i>group by</i> order). Use Sort First By to set the ascending (A to Z) or descending (Z to A) order of the first or most-significant column.
Then By	Use the first Then By sort-order field to set the second sorting order. Use the second Then By sort-order field to set the third sorting order.
More, Less buttons	Click More if you need an additional Then By field. Click Less to remove one.

Saving Copies of Active Channels and Filters

You can save copies of active channels or their filters to modify them later. This is useful to retain an original channel or filter as is, but use a copy of it for a new resource.

To save a copy of an active channel under a new name:

1. Right-click the filter name in the header, and choose **Save Active Channel As**.

This opens the Active Channels Selector dialog, which shows the Active Channels resource

tree.

2. Navigate to where you want to save the channel, enter a new name for it, and click **OK**.

Use a copy of the filter for an active channel independently, or as a basis for other filters. Right-click the filter name in the header, and choose **Save Filter**. This opens the Filter Selector dialog, which shows the Filters resource tree. Navigate to where you want to save the filter, enter a new name for it, and click **OK**.

Best Practices to Optimize Channel Performance

This topic compares active channels, reports, query viewers, and trends in terms of goals and optimal resources for various use cases.

Active Channels or Reports?

Active channels are the better choice if you would rather see results streaming in as the queries proceed, rather than wait for the results to appear in one view in a report.

However, if speed of results is your goal, you might want to run a report instead. The total completion time of an equivalent report would be faster than the total time it takes for the channel to load 100%. This is because the active channel runs multiple smaller queries instead of one large query to display initial results quicker.

See also ["Queries" on page 319](#), ["Building Trends" on page 467](#), and ["Understanding the Reporting Workflow" on page 388](#); and "Query and Trend Performance Tuning" in the [ESM Administrator's Guide](#).

Active Channels or Query Viewers?

[Query Viewers](#) behave more like reports in that they issue a single query against the database and return all results in one batch instead of the streaming progression of results from an active channel. Query viewers are most suitable if you have to slice and dice these query results further, for example, by changing the sort columns, changing types of charts/grid, and so on. These operations are performed on the client side with the results of the already-executed query. If you were using active channels instead, these types of changes would result in a re-running the query as many times as you sort columns.

See also ["Query Viewers" on page 340](#).

Active Channel Query Time Ranges

Take note of the query time range in your active channels. The more hours you are querying, the slower the results are to load. An active channel shows results in minutes if you are

querying a few hours of data. But the channel might start taking several hours to query larger time ranges that span more than 24 hours of data.

If you are querying over more than a day's worth of data, we recommend running a report (using queries and trends) or a query viewer instead of using an active channel.

Active Channel Filters

The more filter conditions you define in an active channel, the more work the channel has to do in the database to evaluate the conditions. A channel that does not have any filter conditions loads data fastest. (This does not mean that the query will run on all events in the database. Only a subset of events are queried, based upon the page you are looking at in the channel.)

If you must use a filter, try to make the filter as restrictive as possible.

Filtering on Indexed Fields

In the ESM, all fields are indexed. There may be search features that can only search a subset of fields, such as query viewer or active channel events, but there are no restrictions based on indexing, since all fields are indexed.

Filtering on Join Fields

The ESM event schema consists of a main **arc_event** table and several other tables. These other tables hold fields related to Annotation, Device fields, Agent fields, Resource References, and so on. If your query has a filter condition on a join field, the resulting channel would have to do more work to evaluate the field.

Continuously Updating Time Parameters

A channel that is “live” (querying against a moving time window and continuously updating the query time ranges) has to do more work than a channel based on fixed time windows. Performance will be better and faster on a channel with a fixed time window than on a live channel. See also [Use of the “Live” Channel from Standard Content](#) for a similar example.

Sorting by End Time or Manager Receipt Time

- By default, all fields can be sorted on the active channel. This sometimes causes performance issues. You can set a property on the Console so that sorting is restricted to the End Time or Manager Receipt Time. If you want sorting on those fields only, follow these steps.

- a. If you have active channels that already sort on fields other than ET or MRT, remove sorting on those other fields manually from the channels. Remove sorting by editing the channel, going to the Sort Fields tab, and change the sorting to either End Time or Manager Receipt Time.
 - b. Stop the Console.
 - c. Add this property to the console.properties file in the system where the Console is installed (for example, C:\arcsight\console\current):

```
console.ui.channel.disable.sorting=true
```

For more information about ESM properties files, see "Managing and Changing Properties File Settings" in the [ESM Administrator's Guide](#).
 - d. Restart the Console.
- Avoid creating channels that are based on one time field but sorted on a different time field. A common cause of poor channel performance is user-created channels with this configuration; that is, a channel based on End Time, but sorted on Manager Receipt Time (or the reverse).



Note: When an event arrives at the Manager so late, that is, beyond its retention period, the Manager adjusts the event's time range so that the event is persisted in the second-to-oldest retained partition. The event is stored in the second-to-oldest partition because the oldest partition may be purged or archived anytime (such as during the data transfer).

This behavior changes the event start and end times, which could cause correlation issues, but the chances of an event being delayed for longer than the retention period are low.

Sorting in Active Channels

By default, the channel has a sort order based upon the time field that was used for creating the channel (End Time or Manager Receipt Time). Note that the sorting operation is done in the database query, so every time you change sort by on any column in your open active channel, ESM has to re-create the complete channel.

The sorting operation can be resource-intensive, especially when millions of events match your filter conditions. Avoid sorting if your filter conditions are not restrictive. For example, the base channel with no filter conditions is normally fastest to load, but it would become the slowest to load if you change its default time based sort order.



Tip: You can use a query viewer instead, that does sorting on the client side with the data it has already queried.

Use of the “Live” Channel from Standard Content

If you are using /All Active Channels/ArcSight System/Core/Live or any similar channel, be aware that the performance of that channel is slower because it has several complex joins (Joins with Annotations, Resource Reference, Device), and performs additional bit-wise operations to evaluate its filter conditions. Depending upon your specific use-case, you can simplify and create your own “Live” channel that is more efficient.

Case Sensitive or Case-Insensitive Conditions?

Wherever possible, use case-sensitive conditions. That will save the extra computation needed for TOUPPER operation required for case-insensitive matches.

I/O Subsystem Performance

Channel query performance is typically limited by the performance of the I/O subsystem on the database. The more events you are inserting, the more load it would cause on the I/O. SAN performance, RAID levels, I/O caches, and so on play a role in how much performance we can obtain.

Diagnostics: Start with Basic Channel Characteristics

To diagnose channel performance issues, start with the most basic active channel to see whether it meets your performance needs, and then keep refining/expanding to come to a point where you can tell what change is affecting performance. Start with the most basic active channel that has the following characteristics:

- Based upon End Time
- No filter conditions (also, make sure to run as an administrator user so that there is no access control filter)
- Query time is two hours ago to Now
- No continuous updates of time parameters

With the above basic active channel, you should see less than a minute wait in starting the channel and doing random scrolls in the channel.

See also:

- ["Status Monitor Events" on page 860](#)
- ["Active Channel Statistics" on page 860](#)

Customizing Columns

You can create active channel columns with customized cell content and presentation formats, tool tip contents, and right-click pop-up values.

You make these changes through the Custom Columns Editor. In the Editor you create new named columns. For each column you select event data fields to display, and if you wish, the HTML formatting to use in its cells. The tool tip option specifies the formatting and content of the tool tips you see when you hover the pointer over cells in that column. The right-click field option sets the event data field to use in columns where there are right-click commands that use field names as arguments such as `Analyze in Channel ...`.

Creating a Custom Column

Where: **Navigator > Resources > Active Channels**


1. Right-click a column header in an active channel and select **Columns > Edit Custom Columns**.
2. In the Custom Columns Editor, click **Add** and name your new column, for example, call it **ABC Vendor**.

If you want all Console users to see the new column (not just administrators), select the **Share with all** check box.

 **Tip:** You can also toggle this option on or off later from the Cell Format tab.

Click **OK**.

3. With the new column selected, click **Field Selector** and select the event attributes you want to display in this column. For example, for ABC Vendor, select **Name** and click **OK**.



The format field is automatically populated with the variable value in Velocity Template format. If appropriate, apply Java-compatible HTML formatting around the field strings. Remember to bracket such formatting with the HTML tag, such as `<HTML><B>${type}</B></HTML>`.

4. Click **Preview** to see how the contents of the **Format** box will look in the active channel.

5. Click the **ToolTip Format** tab to define a tool tip.
6. Select a target event attribute in the **Right-Click Field** menu to populate variable right-click commands, when applicable.
7. Click **Rename** or **Remove** to change or remove selected items in the **Custom Columns** list.
8. Click **Apply** to put your changes into effect and **Close** to close the Custom Columns Editor.

You can edit custom columns after they are created, including toggling on/off the “Share with all” settings for a column, renaming it, changing its Field Selector mappings, and so forth.



Note: Custom columns are not available as arguments for inline filtering.

Showing a Custom Column

A new custom column is immediately available for use in the Console. Right-click the column header in an active channel and choose **Customize Columns >Add Column** to add the new column to the active channel. Custom columns show up in the Available Fields list under Custom Column. If a column is configured as Share with all, it is available to all administrators. If not, it is available only to the user who created it. For more information, see ["Adding, Replacing, or Removing a Column" on page 239](#).

Advanced Example: Creating a Custom Column with Velocity Template

Custom columns can display different contents based on external conditions. Use the Velocity template language to specify these conditions.

To create a custom column that displays a particular image when an event's target is in a specific Zone, create the custom column as described previously, but specify Velocity template-language script in place of the HTML format.

The code in the **Format** text box might look like this:

```
<HTML>
#if (($targetZoneUri.length()>0) &&
    ($targetZoneUri.startsWith("/All Zones/
    System Zones/Public Address Space Zones/
    Ford Motor Company")))
    <IMG src="file:///c:/fordlogo.gif" />
#end
</HTML>
```

Using Dashboards

Dashboards are a graphical display of data gathered from one or more:

- [Data Monitors](#)
- [Query Viewers](#)

Dashboards can display data in a number of graphical formats, including pie and bar charts, tables, and custom layouts.

Administrators can control visibility of, or access to, dashboards, query viewers, and data monitors by changing access control lists (ACLs) as needed. For more information on general use of ACLs on any resource, see ["Managing Permissions" on page 107](#).

With ACLs, administrators can also control which users are allowed to *deploy* (enable) or *un-deploy* (disable) a data monitor.

Monitoring Dashboards

You can organize and present events displayed by data monitors and query viewers on the dashboard. Basic tasks include loading and displaying dashboards; inspecting events; using zoom, slide show, or manipulating the views in various ways; working with dashboard layouts; saving dashboards, and so on.

Where: Navigator > Resources > Dashboards

To load dashboards:

1. Select **Views > Show Dashboard** to open the Load Dashboard dialog box.
2. Expand the dashboard groups to locate the dashboards you want to include in your display.
3. Select the checkboxes next to the dashboards you want to include.
4. Click **OK**.

To display a dashboard:

Right-click a dashboard and select **Show Dashboard**.

To inspect events in dashboards:

1. Open a dashboard containing events.
2. Select the events, right-click, and select **Show Event Details** for Last *N* Events data monitors or **Show details** for other types of data monitors.
 - If you select events from a Last *N* Events data monitor, the details appear in the Event Inspector.
 - If you select events from any other data monitor or query viewer, a new view opens in the Viewer panel for you to investigate.

You can drill down on grid, graph, or chart views.




Tip: By default on a data monitor, the displayed channel uses the same columns as the default Standard Field Set.

If a custom field set is defined for the data monitor **Select Field Set** option, the drill-down channel will use that field set. (See ["Data Monitors" on page 959](#) for information on creating data monitors and defining settings for them.)

You can add or remove columns in the active channel. To do so, right-click on the active channel column headers to get the **Customize Columns** option.

For example, to investigate a data monitor pie chart display, either double-click the chart, or right-click and select **Analyze in Channel > Create Channel** and choose a create channel option, as shown (the menu option for a pie chart display of a query viewer may slightly vary with an additional menu level).



An active channel is displayed showing more detail about the events or resources in the original data monitor or query viewer display. If the channel came from a data monitor, the channel uses the field set columns defined for use in the data monitor **Select Field Set** option. If no field set is defined, the data monitor uses standard field set columns.

To edit dashboard elements:

Right-click in the element and select **<Dashboard element>Edit**.

See also:

- [Editing a Data Monitor](#) and [Moving or Copying a Data Monitor](#)
- ["Managing Query Viewers" on page 341](#)

To drill down to other resources:

If your dashboard contains data monitors and query viewers that have drilldowns, you can view these drilldowns by right-clicking the particular dashboard element (data monitor or query viewer) and selecting **Drilldown > [drilldown name]** from the context menu. The Console then displays the selected element's edit panel at the Drilldowns tab.

If the selected dashboard element does not have a drilldown, you are presented with the option to create one for that element. See ["Adding a Drilldown" on page 270](#).

To display dashboards in a slideshow rotation:

1. In the Viewer panel, select **Views > Slideshow > Interval** in the Console's menu.
2. Use **Interval** to set the number of seconds to pause on each dashboard.
3. Select **Views > Slideshow > Start**, or use the toolbar button , to begin the slide show.

Slide shows appear full-window. Also, **Tile Best Fit** is the best display choice in slideshow dashboards so all data monitors are visible. Use **Views > Slideshow > Stop**, or the toolbar button, to end a slideshow and return to the previous view.

To rearrange elements in dashboard layouts:

Drag and drop the elements on the desired location in the dashboard. You can also click an element's header and drag it to another location.

To use dashboard right-click menu options:

Right-click a dashboard element in a dashboard to use the **Dashboard** subcommands on its context menu. The nature of the element (data from a data monitor or query viewer) determines which commands are applicable and enabled.

See ["Dashboard Context Menu Commands" on page 910](#) for descriptions of options.

To zoom in and out of dashboards:

Right-click a dashboard element and choose **Dashboard>Zoom In** or **Dashboard>Zoom Out**.

To fit dashboard elements:

1. Right-click a dashboard element.
2. Select **Dashboard>Fit in Dashboard**.

To change dashboard layouts:

1. Click the **Layout** button at the lower-right corner of the dashboard in the Viewer panel
2. Select a tab or tile option.
3. Right-click and select **Save Dashboard**.

To close a dashboard:

In a dashboard, right-click and select **Close Dashboard**.

Creating or Editing a Dashboard

When you create a dashboard, the ability to add data monitors is automatically available.

Where: Navigator > Resources > Dashboards

To create or edit a dashboard:

1. If you are creating a dashboard, right-click a dashboard group and select **New Dashboard**.
If you are editing a dashboard, right-click an existing dashboard and select **Edit Dashboard**.
If you selected New Dashboard, an untitled dashboard appears in the Viewer panel and the Data Monitors tab automatically comes forward so you can choose monitors to add.
2. Provide a name if this is a new dashboard.
3. If you want to add data monitors now:
 - a. On the **Data Monitors** tab, navigate through the groups of existing data monitors to find ones you want to add to the dashboard.
 - b. Select a data monitor to add, right-click it and select **Add to Dashboard As**, then select an applicable display format:

Display Options for Dashboards

Display Format	Description
Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. Applies to data monitors and query viewers.
Bar Chart Table	A grid of proportional bar elements. Applies to data monitors.
Horizontal Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. This format forces the bars to run left-to-right rather than up-and-down. Applies to data monitors and query viewers.
Pie Chart	Shows data as a circle with proportional wedges for elements. Applies to data monitors and query viewers.
Statistics Chart	Displays Moving Average data monitors, especially those that contain and need to arrange (overlay) multiple graphs in one monitor space. Compare Statistics Chart to Tile, which arranges individual-graph monitors into fixed arrays. Applies to data monitors.
Table	Displays data as a grid. Applies to data monitors and query viewers.
3D Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. The graph also has a third axis (depth) to display more data and can be rotated by dragging. Applies to data monitors.
Stacked Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. Applies to query viewers.
Tile	Arranges individual Moving Average data graphs into separate, fixed positions on a data monitor, when multiple graphs are present. Compare Tile to Statistics Chart, which displays multiple graphs (overlaid) in the same monitor space. Applies to data monitors.

- c. Repeat the above step to add other data monitors, as needed. When you've finished, right-click the dashboard in the Viewer panel and choose **Save Dashboard**.
4. Additional tasks:
 - To add query viewers, see ["Adding a Query Viewer to a Dashboard" on page 258](#).
 - To add data monitors, see ["Adding a Data Monitor to a Dashboard" on the next page](#).
5. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
6. Click **OK**.

Next steps:

Add elements to this dashboard:

- ["Adding a Data Monitor to a Dashboard" below](#)
- ["Adding a Query Viewer to a Dashboard" on the next page](#)

Adding a Data Monitor to a Dashboard

Where: Navigator > Resources > Dashboards

To add a data monitor to a dashboard:

1. Right-click a dashboard and select **Show Dashboard**.
2. On the Data Monitors tab, navigate through the groups of existing data monitors to find ones you want to add to the dashboard.
3. Right-click a data monitor and select **Add to Dashboard As**, then select an applicable display format:

Display Options for Dashboards

Display Format	Description
Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. Applies to data monitors and query viewers.
Bar Chart Table	A grid of proportional bar elements. Applies to data monitors.
Horizontal Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. This format forces the bars to run left-to-right rather than up-and-down. Applies to data monitors and query viewers.
Pie Chart	Shows data as a circle with proportional wedges for elements. Applies to data monitors and query viewers.
Statistics Chart	Displays Moving Average data monitors, especially those that contain and need to arrange (overlay) multiple graphs in one monitor space. Compare Statistics Chart to Tile, which arranges individual-graph monitors into fixed arrays. Applies to data monitors.
Table	Displays data as a grid. Applies to data monitors and query viewers.

Display Options for Dashboards, continued

Display Format	Description
3D Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. The graph also has a third axis (depth) to display more data and can be rotated by dragging. Applies to data monitors.
Stacked Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. Applies to query viewers.
Tile	Arranges individual Moving Average data graphs into separate, fixed positions on a data monitor, when multiple graphs are present. Compare Tile to Statistics Chart, which displays multiple graphs (overlaid) in the same monitor space. Applies to data monitors.

4. To save the dashboard, right-click it and choose **Save Dashboard**. If this is a new dashboard, navigate to the group where you want to save the dashboard, enter a name for the new dashboard, and click **OK**.

Adding a Query Viewer to a Dashboard

You can add a query viewer result to a dashboard as follows:

Where: Navigator > Resources > Query Viewers

1. If you have identified an existing dashboard to which you want to add the query viewer, open the dashboard in the viewer and make sure it is the focus. If you want to add the query viewer to a new dashboard, continue to the next step.
2. Right-click a query viewer and select **Add to Dashboard As >**, then select an applicable display format (see ["Dashboard Display Formats" on the next page](#)).

The query viewer result is displayed on the open dashboard. If a dashboard is not displayed, a new untitled dashboard is created for the query viewer result.

3. Save the existing dashboard.

Or if this is a new dashboard:

- a. Right-click the title bar of the dashboard and choose **Save Dashboard As**.
- b. In the popup dialog, navigate to the group where you want to save the dashboard, enter a name for the dashboard, and click **OK**.

By default, this new dashboard is a regular dashboard. If you want to change it to a custom view dashboard, see ["Using Custom View Dashboards" on page 261](#).

You can add multiple query viewer results sets along with other resources to a single dashboard.

Dashboard Display Formats

The available display options depend on the nature of the dashboard element.

Display Options for Dashboards

Display Format	Description
Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. Applies to data monitors and query viewers.
Bar Chart Table	A grid of proportional bar elements. Applies to data monitors.
Horizontal Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. This format forces the bars to run left-to-right rather than up-and-down. Applies to data monitors and query viewers.
Pie Chart	Shows data as a circle with proportional wedges for elements. Applies to data monitors and query viewers.
Statistics Chart	Displays Moving Average data monitors, especially those that contain and need to arrange (overlay) multiple graphs in one monitor space. Compare Statistics Chart to Tile, which arranges individual-graph monitors into fixed arrays. Applies to data monitors.
Table	Displays data as a grid. Applies to data monitors and query viewers.
3D Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. The graph also has a third axis (depth) to display more data and can be rotated by dragging. Applies to data monitors.
Stacked Bar Chart	Shows data as a series of proportional bar elements and may include bar segmentation to subdivide the data. Applies to query viewers.
Tile	Arranges individual Moving Average data graphs into separate, fixed positions on a data monitor, when multiple graphs are present. Compare Tile to Statistics Chart, which displays multiple graphs (overlaid) in the same monitor space. Applies to data monitors.

Managing Dashboard Groups

The groups in the Dashboard tab of the Navigator panel's Dashboard resource tree store individual dashboards or other dashboard groups. You use groups within groups to help organize larger numbers of resources.

You can manage groups by drag-and-drop. You can move or copy dashboards or groups within the Dashboards resource tree. And deleting a group also deletes the resources it contained.



Note: To copy multiple resources at once, use Copy and Paste. You can drag and drop only one resource at a time.

Where: Navigator > Resources > Dashboards

To create a dashboard group:

1. Right-click a group and choose **New Group**.
2. Enter a name in the group's text field.
3. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
4. Press **Enter**.

To rename a dashboard group:

1. Right-click a group and choose **Rename**.
2. Enter a name in the group's text field.
3. Press **Enter**.

To edit a dashboard group:

1. Right-click a group and choose **Edit Group**.
2. In the Group Editor, edit the **Name** and **Description** text fields.
3. Click **OK**.

To move or copy a dashboard group:

1. Navigate to a group and drag it into another group.
2. Choose **Move** to move the group, **Copy** to make a separate copy of the group, or **Link** to create a copy of the group that is linked to the original group.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected

when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete a dashboard group:

1. Right-click a group and choose **Delete Group**.
2. In the dialog box, click **Yes**.

Using Custom View Dashboards

You can create custom layouts of dashboard data. Also known as *image dashboards*, custom view dashboards enable you to create custom views of dashboard data from query viewers and from data monitors over an imported image, such as a geographical map.

From the Console, custom view dashboards are displayed in the Console user's preferred Web browser.



Note: Refer to the Support Matrix document applicable to your ESM version for the official list of supported Web browsers.

In custom view dashboards, only drill-downs to other dashboards are supported. Custom view dashboards refresh event data at the same rate as regular dashboards.

Displaying Custom View Dashboards

There are several ways to switch from the regular dashboard view to the custom view dashboard view. Each of these methods loads the custom view dashboard editor in the Viewer Panel in the custom view dashboard's *View* mode with the last configuration you saved.

From the regular dashboard view in the Viewer panel:

1. In the Navigator panel, go to Dashboards.
2. Double-click the dashboard you want to open, or right-click it and select **Show Dashboard**.
3. Do one of the following:
 - Click the Layout Selector button at the bottom of the display () and select **Custom Layout**.
 - Right-click the dashboard tab and select **Choose Layout > Custom Layout**.

- Right-click any data monitor in the dashboard and select **Dashboard > Choose Layout > Custom Layout**.

The dashboard is displayed on your preferred Web browser.



Note: Important reminders about custom view dashboards:

- Save a new or edited dashboard before accessing the custom view dashboard.
If you are creating a new dashboard or have edited an existing one to use the Custom Layout, first save the dashboard to establish the custom layout for this dashboard on the Manager.
- Custom view dashboards are displayed with the default chart and color settings.
User customizations to chart settings and color selections applied to dashboards in the regular viewer are not applied to the custom view dashboard view.
- Custom view dashboard backgrounds scale to fit the available Viewer panel space.

For more about selecting and working with background images for custom view dashboards, see ["Loading a Background Image" on page 264](#).

To refreshing the custom view dashboard layout:

On the custom view dashboard, click the **Reload** button () on the top right of the menu bar.

Reverting to the Regular Dashboard View

You cannot display the regular view and the custom view of the same dashboard at the same time. Therefore, every time you show a dashboard on the Viewer panel and then display its custom view, the dashboard's tab on the Viewer panel becomes empty.

If your dashboard is set to display in a layout other than Custom Layout (for example, **Tab**), and if you displayed the dashboard in custom view, that dashboard is always displayed in custom view until you revert to regular view, as described below:

1. Close the browser where the custom view dashboard is displayed.
2. Right-click the dashboard's tab in the Viewer panel and select **Close** to close the empty dashboard.
3. Double-click the dashboard in the Navigator panel, or right-click it and select **Show Dashboard** to re-open the dashboard in regular view.



Note: If you edited a dashboard and changed its Layout setting to **Custom Layout**, then the dashboard is always displayed in custom view when you right-click on it and select **Show Dashboard**. To revert to regular view, edit the dashboard again and change its Layout setting to one of the other options.

Working with Custom View Dashboards

Custom view dashboards provide two modes: *View* mode for monitoring and investigating events, and *Edit* mode, for arranging and customizing background elements.

Custom view dashboards open in the *View* mode. In *View* mode, you can select a dashboard element then choose a different display format.

If there is no background associated with a dashboard from when it was created using the Dashboard editor, the dashboard is displayed with a white background; otherwise you see the last background that was added. The dashboard elements are rendered in evenly distributed rows.

To select View mode:

1. In the custom view dashboard top menu bar, select **View**.
2. From the drop-down menu, select the dashboard element and the desired display format.
Some dashboard elements may not offer other display format options.

Arranging Custom View Dashboards

In *Arrange* mode, you can customize the dashboard layout, toggle data monitors on and off, and upload a background image.

When you switch to Arrange mode, chart-type data monitors appear with a yellow background. You can relocate, resize, and reshape all types of data monitors anywhere in the custom view dashboard view.



Note: Changes saved to a custom view dashboard refreshes the dashboard on all ArcSight Consoles attached to the Manager.

If the ArcSight Manager supports more than one ArcSight Console, any custom view dashboard changes saved on one Console will refresh that dashboard on the other Consoles attached to the Manager.

To select Arrange mode:

In the custom view dashboard's top menu bar, select **Edit > Arrange**.

Loading a Background Image

You can upload a background image to the custom view dashboard. The image you select will be stretched to fit the available display space in the Viewer panel, so for best results, select an image with adequate size and proportion to fill the space.

1. In the Navigator panel, go to **Files**.
2. Right-click a folder (for example, your personal folder) and select **New File**. In the Inspect/Edit panel, give the file a name, and click **Upload**.
3. In the popup, browse to the location of the image file and click **OK**. The file is stored as a resource.
4. In the Navigator panel, go to **Dashboards**. Double-click the dashboard you want to open in the Viewer panel, or right-click it and select **Show Dashboard**.
5. Click the Layout Selector button at the bottom of the display () and select **Custom Layout** to open the custom view dashboard on the browser.
6. On the custom view dashboard, select **Edit > Upload New Background Image**.
7. In the Upload New Background Image popup, browse to the location of the image and select it. Then click **Upload**.



Tip: If the background image is not displayed right away, refresh the custom view dashboard.

To refresh the Custom Layout view, click the **Reload** button ( **Reload**) at the top right of the toolbar.

Selecting a Previously Uploaded Background Image

If you have previously uploaded a background image that you want to load as the custom view dashboard background, or you want to use an image, use the File menu resource. This assumes you have previously uploaded an image using the File resource.

Where: Navigator > Resources > Dashboards

1. Double-click the dashboard you want to open in the Viewer panel, or right-click it and select **Show Dashboard**.
2. Click the Layout Selector button at the bottom of the display () and select **Custom Layout** to open the custom view dashboard on the browser.
3. On the custom view dashboard, select **Edit > Choose an Existing Background Image**.
4. In the Select a File popup, select the image you uploaded to the Files resource. Then click **OK**.
5. Click **Save** on the top right of the custom view dashboard's menu bar.

Verifying the Background Image

You can verify that a background image has been attached to a custom view dashboard without launching the custom view dashboard itself.

Where: Navigator > Resources > Dashboards

1. Right-click the dashboard and select **Edit Dashboard**.
2. In the Inspect/Edit panel, verify the values in the following fields:
 - Layout = Custom Layout
 - Background = URI to the file resource for the image, or path to the image file in another drive

Removing a Background Image

You can remove the background image directly from the custom view dashboard or from the dashboard's Inspect/Edit panel.

To remove a background image from the custom view dashboard:

1. Display the dashboard in custom view.
2. Click **Edit** from the top menu bar and select **Clear Background**.
3. Click **Save** on the top right of the custom view dashboard's toolbar.

Custom View Dashboard Context Menu Options

Both the View and Arrange modes offer the following context menu options (**Ctrl**, **Alt**, or **Shift** + left click) for custom view dashboards:

Context Menu Options for Custom View Dashboards

Option	Description
Save	Save the current dashboard layout. This option becomes available when you have selected a different layout using the View As context menu option.
Data Monitor: Enable/Disable	Enable or disable the data monitor. For more about enabling and disabling data monitors, see "Enabling or Disabling a Data Monitor" on page 278 .

Context Menu Options for Custom View Dashboards, continued

Option	Description
View As	Change the data monitor view to a graphical format supported for the data monitor type. For more about data monitor views, see "Dashboard Display Formats" on page 259 .
Auto Arrange	Let the Console rearrange the dashboard elements in the best possible fit, horizontally from left to right. Any remaining available space is moved to the bottom.
Choose Colors	Set foreground and background colors. If the display format is a table, you can additionally associate row or cell colors with returned values.

Using Data Monitors

You populate dashboards with data monitors, which you most often select from the Data Monitors resource tree in the Navigator panel (under Dashboards). Pre-defined data monitors are provided. You can create, edit, and delete your own data monitors.

Administrators can limit visibility of, or control access to, data monitors by changing access control lists (ACLs) as needed. For more information on general use of ACLs on any resource, see ["Managing Permissions" on page 107](#).

With ACLs, administrators can also control which users are allowed to *deploy* (enable) or *un-deploy* (disable) a data monitor.

Creating a Data Monitor

Where: Navigator > Resources > Dashboards > Data Monitors tab

1. Right-click a data monitor group and choose **New Data Monitor**.
2. In the Data Monitor Editor, select a **Data Monitor Type** from the drop-down menu.

Data Monitor Types

Data Monitor Type	Description
"Asset Category Count Data Monitor" on page 960	Enumerate the number of real-time hits (events) that occur per asset category, by priority, within a time interval.
"Event Correlation Data Monitor" on page 961	Provide flow-volume level correlation between two different event streams (based on two different specified filters).
"Event Graph Data Monitor" on page 964	Draw real-time diagrams of selected event activity. Automates the graphing of attacks in real-time. The <i>manual</i> operations are described in "Graphing Attacks" on page 287 .
"Geographic Event Graph Data Monitor" on page 965	Draw a real-time geographic map of selected events. In effect, it does automatically and in real-time what you can do manually, as described in "Graphing Attacks" on page 287 .
"Hierarchy Map Data Monitor" on page 966	Draw an image made up of proportionally sized panels where each panel represents a group of events selected by group fields selected in the source node identifier. A source-node criteria could be a combination of fields. The Hierarchy Map data monitor includes several enhancements, as described in "Hierarchy Map Features" on page 967.
"Hourly Counts Data Monitor" on page 976	Display the total count of events on an hourly basis along with their Priority.
"Last N Events Data Monitor" on page 977	Order events based on a specified configuration. In the Table Viewer, the monitor displays the most recent events by Priority, Event Name, Protocol, and Category. With the BarChartTable configuration, the order is by Priority and Event Name. The PieChart configuration is ordered by Priority.
"Last State Data Monitor" on page 978	Provide an extra level of abstraction that you can use to simplify the information presented to operators. Sometimes called <i>indicator lights</i> or <i>heads-up displays</i> , these monitors show graphics that translate more complex values into simple, rapidly observable results such as green/amber/red signal lights or checkmark/asterisk/exclamation point symbols. Last State data monitors could also be called <i>most recently known state</i> monitors.

Data Monitor Types, continued

Data Monitor Type	Description
"Moving Average Data Monitor" on page 982	Display the moving average of events by a selected data field. The display provides a running count of events within a specified time frame and generates an event when the moving average changes significantly.
"Rules Partial Match Data Monitor" on page 985	Display rules that have partial matches and the total number of partial match events within a specified time frame. For more information on partial matches, see "Managing Rule Actions" on page 532 .
"Statistics Data Monitor" on page 986	Provide a broader generalization of Moving Average data monitor functionality, except that it allows selection of other statistical methods in addition to Moving Average. Statistical methods include Average, Moving Average, Standard Deviation, Skew and Kurtosis, as well as Moving Average. These added capabilities could be used to detect anomalous behavior that could not be detected using moving average alone.
"System Monitor Data Monitor" on page 988	Provide measurements based on ArcSight Manager internal monitoring system Java classes and attributes. A number of system monitors that might be particularly useful to ArcSight administrators are provided as predefined System Data Monitors that you can include in your dashboard displays to monitor system performance.
"System Monitor Attribute Data Monitor" on page 989	Similar to System Monitor, except that, rather than provide measurements for all attributes of a specified Java class, focus on a single specific attribute of a given ArcSight Java class. Used primarily for measurements on attributes that provide complex data structures.
"Top Value Counts Data Monitor" on page 990	Display top events by selected data field, the total number of events, and the event Severity within the total number of events with the Table and BarChartTable viewer configurations.

- Based on the data monitor type you have selected, specify values and options in the applicable fields to define the data monitor's data collection. Details on fields and appropriate values are given in the information about each data monitor type.



Note: Depending on the permissions associated with the user group to which you belong, you may or may not have an option to **Enable** (*deploy*) or disable (*un-deploy*) the data monitor. For more information, see ["Enabling or Disabling a Data Monitor" on page 278](#).

- If the data monitor uses data fields for evaluation, use the **Variables** tab to create a new specialized field, if necessary.

The following data monitors support variables:

- Event graph
- Hierarchy Map
- Last N Events
- Last State
- Moving Average
- Statistics
- Top Value Counts (bucketized)

If you select a data monitor that does not support variables, the Variables tab is disabled.

You can also add a global variable anywhere fields can be added. See ["Adding a Global Variable to a Data Monitor" on page 620](#).

5. If the Data Monitor type supports drill downs to other resources, use the Drilldown tab to configure it. The following types of Data Monitors support drilldowns:

- Event Graph
- Hierarchy Map
- Last N Events
- Last State
- Moving Average
- Statistics
- Top Value Counts (Bucketized)

See ["Adding a Drilldown" on the next page](#) for instructions.

6. Click **OK**.

To add the new monitor to the current dashboard, right-click it and choose **Add to Dashboard As**.

Editing a Data Monitor

1. Access the data monitor for editing in one of two ways:
 - In Navigator > Resources > Dashboards > Data Monitors tab, right-click a data monitor and select **Edit Data Monitor**.
 - If a dashboard containing the data monitor is already displayed, hover the cursor over that data monitor in the dashboard's Viewer panel, right-click, and select **Data Monitor > Edit**.

2. In the Data Monitor Editor, edit the applicable fields.

See ["Creating a Data Monitor" on page 266](#) and ["Data Monitors" on page 959](#) for field details on all data monitors.

For customized view options on Last State data monitors, see the [Last State Data Monitor](#) topic.

3. Click **OK**.

Deleting a Data Monitor

Where: **Navigator > Resources > Dashboards > Data Monitors tab**

1. Right-click a data monitor and choose **Delete Data Monitor**.
2. In the dialog box, click **Yes**.

Managing Drilldowns from Data Monitors

Drilldowns provide the ability to investigate details about resources related to what is displayed by query viewers or data monitors. You can get more focused views on particular aspects of a single item, such as an asset, case, event, and so on, in the query result.

Adding a Drilldown

You can configure query viewers and data monitors to drill down to one or a combination of the following resources:

- Active channels
- Dashboards
- Query viewers
- Reports

Each drilldown type has its own options. After you have added one or more drilldowns, Console users can select one by right-clicking on the result and selecting **Drilldown > [drilldown name]** from the context menu.



Note: In a Custom View Dashboard and on the ArcSight Command Center, only drilldowns to dashboards are supported.

You *can* create drilldowns from these types of data monitors:

- Event graph
- Hierarchy map

- Last N Events
- Last State
- Moving Average
- Statistics
- Top Value Counts

You *cannot* drill down to resources from the following data monitors:

- Asset Category Count
- Event Correlation
- Geographic Event Graph
- Hourly Counts
- Rules Partial Match
- System Monitor
- System Monitor Attribute

Where: Navigator > Resources > Dashboards > Data Monitors tab > data monitor > Drilldowns tab

To add a drilldown from the data monitor:

1. Access the **Drilldowns** tab in one of two ways:
 - Right-click on the query viewer or data monitor results in a dashboard and select **Drilldowns/Edit Drilldowns** to open the editor to the **Drilldowns** tab.

Or

 - Right-click on a query viewer or data monitor in the Navigator panel and select the **Edit** option, then select the **Drilldowns** tab.
2. Click **Add** ( **Add...**) to open the Add Drilldown panel.
3. In the Destination field, select a resource type, for example, Dashboards.

The screenshot shows the 'Add Drilldown' dialog box. The 'Destination' dropdown menu is open, displaying a list of options: 'Dashboards', 'Active Channels', 'Dashboards', 'Query Viewers', and 'Reports'. The first 'Dashboards' option is highlighted. The dialog also contains a 'Menu Label' text field with the instruction 'Give your drilldown a descriptive name.', a 'Description' text field with the instruction 'Add more descriptive information about this drilldown.', and four buttons at the bottom: 'Next', 'Finish', 'Cancel', and 'Help'.


Then choose the corresponding specific resource, for example, My_Dashboard.

4. Enter a menu label (defaults to the specific resource's name). This label will represent this drilldown when the user right-clicks and selects Drilldowns on the Viewer panel.
5. Enter an optional description containing useful information about the drilldown.
6. Set the remaining options based on your destination resource:


Options for the Drilldown's Resource Destinations

If resource type is ...	Follow these steps ...
Active Channels	For an active channel destination, the settings in the Channel Display Options tab are not required; you may click Finish . If you want to set display options: - Select a field set from the drop-down list and click OK. - Change the Sort By field from the drop-down list and the sort order. - Click Finish.
Dashboards	Click Finish . You are done.

Options for the Drilldown's Resource Destinations, continued

If resource type is ...	Follow these steps ...
Query Viewers	For a query viewer destination, field mapping is required: - On the Field Mapping tab, click Add to display a dropdown list of source fields. You must define at least one field map. The source fields are from the source query viewer (the one you are drilling down <i>from</i>). The mapping condition is always set to =. - Under the Destination Field column, select a field from the destination query viewer (the one you are drilling down <i>to</i>).  The Drilldown definition shown in the example maps the source query viewer/data monitor “Name” column to the target query viewer/data monitor “Name” column. This constructs the following drilldown filter: <pre><target>.Name = <source>.Name</pre> where <source>.Name is replaced by the actual value from the source query viewer/data monitor row. If there are no eligible field mappings, you cannot complete the drilldown definition; the Finish button is disabled. You can add or remove field mappings, but your choices are limited to the columns already provided in the query viewer. - On the Display tab, you can choose to show (check) or hide (uncheck) the data fields in the drilldown result. - On the Sort tab, you can click Add to select the columns to specify the sort order of the resulting values. For each added column, change the sort order to ascending (the default) or descending. - Click Finish.

Options for the Drilldown's Resource Destinations, continued

If resource type is ...	Follow these steps ...
Reports	For a report destination, the settings in the Report Display Options tab are not required. To use the parameters set for the report, click Finish. If you want to change the drilldown's display options: - Click Add to display a list of the destination report's custom parameters, then select a parameter. - Under the Value column, select the field whose value will be used for the parameter.  - Click Finish.


7. Repeat the process to add multiple drilldowns as required.

The drilldowns you added will be available for selection when you view the data monitor or query viewer results. From those resources, the drilldowns are displayed for selection in the order they were created. The first drilldown is automatically the default drilldown of choice.

Tips on drilldown definitions:

- If there is only one drilldown, this is the default drilldown for that resource. If there are multiple drilldowns, the first drilldown is the default. You can change the order on the Drilldowns tab.
- When you run the query viewer results or view a data monitor, right-click, and select **Drilldown**, the selection list displays the list of drilldowns defined for that resource. The default drilldown is at the top of the list, and the remaining drilldowns are displayed in the sequence as they appear on the data monitor or query viewer's Drilldowns tab.
- You can define drilldowns for multiple fields of different data types. For example, you can define a drilldown to return a combination of event name and IP address. The first step would be to define a base query viewer to return these fields in a result, and then, as a next step, add a drilldown and select that query viewer to use as the “Drill down to” query viewer.
- You cannot define drilldowns to go to fields that are SQL functions.

Example of drilldowns added to a query viewer



Editing a Drilldown

Purpose: To change drilldown settings defined for a data monitor

Where: Navigator > Resources > Dashboards > Data Monitors tab > data monitor > Drilldowns tab

To edit a drilldown:

1. Open the editor for the query viewer or data monitor you want to edit.
2. Click the **Drilldowns** tab.
3. Select the drilldown you want to edit and click **Edit**  Edit...

The drilldown dialog for this drilldown is displayed. Change the fields and options as described in ["Adding a Drilldown" on page 356](#).



Note: You can also edit the drill down from the query viewer or data monitor results. Right-click and select **Drilldown > Edit Drilldowns**. Selecting this command opens the editor for the query viewer or data monitor at the Drilldowns tab.

Changing the Default Drilldown

Purpose: To change the top selection of the drilldown's selection list

Where: Navigator > Resources > Dashboards > Data Monitors tab > data monitor > Drilldowns tab

When you run the query viewer results or view a data monitor, right-click, and select **Drilldown**, the selection list displays the list of drilldowns defined for that resource. The default drilldown is at the top of the list, and the remaining drilldowns are displayed in the sequence as they appear on the Drilldowns tab. This default position is not affected by any sorting of drilldowns.

To change the default drilldown:

1. Open the editor for the data monitor or query viewer you want to edit.
2. On the Drilldowns tab under the Default column, click the button corresponding to the drilldown you want as the default and save.

The default drilldown will appear at the top of the selection list the next time you right-click on the query viewer results or data monitor and select **Drilldown**.

Sorting or Changing the Order of Drilldowns

Purpose: To change the order of drilldowns on the list

Where: Navigator > Resources > Dashboards > Data Monitors tab > data monitor > Drilldowns tab

If you create multiple drilldowns to different resource types, the Drilldowns tab displays the drilldowns in the sequence they were created. This initial sort order affects the selection list if you right-click the data monitor or query viewer results on the Viewer panel and select **Drilldowns**.

You can re-order the drilldowns in two ways:

- Sorting the drilldowns
- Moving specific drilldowns up or down the list

To change the sort order:

1. Open the editor for the data monitor or query viewer you want to edit.
2. Click the **Drilldowns** tab and click **Sort** () on the toolbar.

Multiple drilldowns on the Drilldowns tab are sorted in two ways, as follows:

- First, the drilldowns are sorted alphabetically according to resource type: active channels, dashboards, query viewers, and reports.
- Next, within the resource type, drilldowns are again sorted alphabetically by their menu labels.

After you click the **Sort** button, clicking it again will not change the sort order.



Note: Even if the default drilldown moves after sorting on the Drilldowns tab, the default will still be at the top of the selection list when you right-click on the data monitor or query viewer results and select **Drilldowns**. If you want to change the default itself, follow instructions in ["Changing the Default Drilldown" on page 360](#).

To move a drilldown's position on the list:

1. Open the editor for the data monitor or query viewer you want to edit.
2. Click the **Drilldowns** tab and select a drilldown. Do not click under the Default column if you are not changing the default drilldown.
3. On the toolbar, click the up  or down  arrow buttons to move the drilldown up or down the list.

Removing a Drilldown

Where: Navigator > Resources > Dashboards > Data Monitors tab > data monitor > Drilldowns tab

You remove any drilldown, including the default drilldown, one at a time. If you delete only the default and you have other drilldowns, the next drilldown on the list becomes the default.

To remove a drilldown:

1. Open the editor for the data monitor or query viewer you want to edit.
2. Click the **Drilldowns** tab.
3. Select the drilldown you want to remove and click **Remove** ( Remove).
4. Repeat as required.

Moving or Copying a Data Monitor

You can move or copy a data monitor as you would any other resource. See "[Moving, Copying, Linking, and Deleting Resources](#)" on page 712.

Where: Navigator > Resources > Dashboards > Data Monitors tab



Note: Regarding data monitors and user permissions

- Users who do not have data monitor deployment permissions can still copy enabled data monitors, but the copies are disabled. Users need both write and deploy permissions to enable or disable a data monitor.
- Users who do not have data monitor deployment permissions can still move data monitors from one group to another if they have:
 - Write permissions on the data monitors they want to move, and
 - Write permissions on the destination group for the move operation

For more about data monitor deployment permissions, see "[Controlling Who Has Permissions to Deploy Data Monitors](#)" on page 116.

Enabling or Disabling a Data Monitor

When a data monitor is enabled (*deployed*) it is actively processing events and updating its display.

When you disable (*un-deploy*) a data monitor, it stops processing events and updating its display. You might choose to disable a data monitor because it is not needed or should not be considered under certain circumstances.

Data monitors can be enabled at time of creation (see ["Creating a Data Monitor" on page 266](#)) or edited later to enable deployment.



Note: Data monitor deployment is controlled through User Access Control Lists (ACLs). Administrators can allow or block users for data monitor deployment permissions.

Depending on the permissions associated with the user group to which you belong, you may or may not have an option to **Enable** (*deploy*) or **Disable** (*un-deploy*) the data monitor.

- Administrators (all users belonging to the Administrators group) have permissions to deploy/un-deploy data monitors.
- To deploy a data monitor, a user needs *both* general data monitor deployment permissions and write permissions to the specific data monitor he or she wants to deploy. Users with permissions to deploy data monitors can deploy only those data monitors for which they have write permissions.
- Administrators can grant permissions to deploy or restrict data monitors to other non-Administrator users through the Access Control Lists (ACLs) editor. For more information, see ["Controlling Who Has Permissions to Deploy Data Monitors" on page 116](#) and ["Granting or Removing Resource Permissions" on page 108](#).

To enable or disable a data monitor from the Editor:



Tip: You can set *operations* permissions on data monitor deployment by editing Access Control Lists (ACLs) on user groups. Administrators can allow or block user groups for data monitor deployment permissions. (This is different than controlling permissions on who has access to the data monitors *resource*.)

To set permissions for *deploying* data monitors, click the **Operations** tab, then click the **Add** button to get the Permissions Selector dialog for operations, select **Deploy** and click **OK**. For more information, see ["Controlling Who Has Permissions to Deploy Data Monitors" on page 116](#).

1. Access the data monitor for editing in one of two ways:
 - In Navigator > Resources > Dashboards > Data Monitors tab, right-click a data monitor and select **Edit Data Monitor**.

- If a dashboard containing the data monitor is already displayed, hover the cursor over that data monitor in the dashboard's Viewer panel, right-click, and select **Data Monitor > Edit**.
2. In the Data Monitor Editor, click the check box for **Enable** to toggle the data monitor on or off.
 - A checkmark indicates the data monitor is enabled (deployed).
 - If the box is unchecked, the data monitor is disabled (undeployed).
 3. Click **Apply** or **OK** on the editor to save your changes.

To enable or disable a data monitor in the Navigator:

You can also enable and disable data monitors in the Navigator by right-clicking data monitors or a data monitor group.

Where: Navigator > Resources > Dashboards > Data Monitors tab

1. Right-click a data monitor or a data monitor group.
2. Choose **Enable Data Monitor** to *deploy* or activate the monitors (if disabled) or **Disable Data Monitor** to *un-deploy* or deactivate (if enabled).

For information about granting permissions to user groups to enable or disable data monitors, see ["Controlling Who Has Permissions to Deploy Data Monitors" on page 116](#).

Overriding a Data Monitor's Last State

Last State data monitors can sometimes display a status that has served its purpose as soon as you have seen it. Once seen, you may want to directly reset or change the status so you can watch for a new status change, without waiting for an automatic system update.

Last State data monitors can be displayed as a ribbon on custom view dashboards. See ["Using Custom View Dashboards" on page 261](#) for more information.

When you see a status in a Last State data monitor that you want to reset, de-escalate, or otherwise override, right-click a cell in the monitor and choose **Override Status**. In the Select dialog box, select the new status and click **OK**.

Managing Data Monitor Groups

Data monitor groups store similar data monitors in a single location. You can create groups within groups to meet enterprise needs.

You can manage one group at a time by drag-and-drop. You can move or copy dashboards or groups within the Dashboards resource tree. Deleting a group deletes the resources it contained.



Note: To copy multiple resources at once, use Copy and Paste. You can drag and drop only one resource at a time.

Where: Resources > Dashboards > Data Monitors tab

To create a data monitor group:

1. Right-click a group and choose **New Group**.
2. Enter a name in the text field.
3. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
4. Press **Enter**.

To rename a data monitor group:

1. Right-click a group and select **Rename**.
2. Enter a new name in the group's text field.
3. Press **Enter**.

To edit a data monitor group:

1. Right-click a group and choose **Edit Group**.
2. In the Group Editor, edit the **Name** and **Description** text fields.
3. Click **OK**.

To move or copy a data monitor group:

1. Navigate to a group and drag it into another group.
2. Choose **Move** to move the group, **Copy** to make a separate copy of the group, or **Link** to create a copy of the group that is linked to the original group.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete a data monitor group:

1. Right-click a group and choose **Delete Group**.
2. In the dialog box, click **Yes**.

To enable or disable data monitor groups:

Data monitors are enabled by default. When you disable data monitors they stop processing events and updating their displays. You might choose to disable a data monitor group because it is not needed or should not be considered under certain circumstances.


You can also enable and disable data monitors individually in the Data Monitor resource tree or Data Monitor Editor.

1. Right-click a data monitor group.
2. Select **Enable Data Monitor** to activate all the monitors in the group (if they are disabled) or **Disable Data Monitor** to deactivate them (if they are enabled).

Optimizing the Evaluation of Event Filters for Data Monitors

The information described here applies to correlation data monitors that use event filters. This topic is written for advanced content authors. The topic describes how to automate the optimization of event conditions to reduce impact on CPU usage.

Evaluating event conditions is one of the most resource-intensive operation in event processing. ESM evaluates event conditions in the sequence they appear on the filter's Filter tab. The following example shows event conditions:



The outer and inner nodes indicate nested conditions. As a guideline, order these conditions from the most economical to the most expensive by putting the least costly condition on top. However, this may not be enough; you must also consider the TRUE/FALSE rate and whether you are using OR or AND operators.

Related topics:

- ["Requirement" below](#)
- ["Automating the Optimization of Filter Conditions" below](#)
- ["Tracing the Optimization" on the next page](#)
- ["Disabling the Optimization Feature" on page 284](#)

Requirement

To optimize the evaluation of event conditions, data monitors must be of the event or correlation type and must be using a filter, because event conditions are defined in the filter. Filters are used not only in data monitors but also in other resources, for example, active channels and query viewers. However, event conditions in filters are only evaluated and optimized in the context of the data monitors that use these filters. The filters themselves are not altered by the evaluation.

Refer to ["Creating or Editing a Filter" on page 306](#) for more information.



Note: A data monitor can use one or more filters. In this case, all filters being used by the data monitor will be optimized in memory.

Automating the Optimization of Filter Conditions

ESM has the ability to evaluate event conditions and if necessary, re-sequence the conditions to be evaluated through this property setting in the `server.properties` file:

```
rule.dm.optimize.evaluation=true
```

This setting automates the process of evaluating and optimizing all of your deployed rules and data monitors. The optimization process occurs in memory and will leave your resources intact. If there are changes, for example, real-time rules are updated, or the user de-activates and then re-activates the rule, optimization runs again.

The following line in `server.log` denotes the end of optimization:

```
End of conditions optimization in <n> events
```



Note: Refer to the *ESM Administrator's Guide's* Configuration section for more information:

- For detailed instructions on how to add settings to the `server.properties` file, refer to the topic, "Managing and Changing Properties File Settings."
- For information on the `server.log` file, refer to the topic, "Configuring Manager Logging."

Tracing the Optimization

The tracing feature described here is optional. However, if you want to use this feature, you must turn off the condition optimization setting. The tracing feature enables you to capture information about how your rules were optimized. The information is stored in `server.log`.

To use the tracing feature, add this property setting in the `server.properties` file:

```
rule.dm.trace.optimize.evaluation=true
```

This setting records in the `server.log` file the original sequence of conditions and how these conditions are re-ordered (optimized).

To locate the information in the log file, search for the the data monitor's name. The following example shows the log statements that include the evaluated data monitor's URI and a sampling of the profiled values. The profile describes the average time in nanoseconds a specific condition took to process. The condition's position is indicated by an index starting with position 0. The profile in the log indicates that the costliest condition at Index 0 took an average of 1,227 nanoseconds over a sampling of 22:

```
[2014-08-01 12:22:19,416] [INFO  
][default.com.arcsight.common.simplerulesengine.engine.b]  
Optimizing /All Data Monitors/ArcSight Foundation/Intrusion Monitoring/Detail/Attack  
Monitoring/Attack Rates/Attack Rates by Service and Zones/Attacker Zones by Service  
Profiled values are  
[ Index=1, avgTime=1185 ,numSamples=22, Index=2, avgTime=1189 ,numSamples=23,  
Index=0, avgTime=1227 ,numSamples=22]
```

The following lines indicate the order of conditions (indicated as nodes) as originally defined. Red numbers correspond to the index:

```
Original nodes are  
Node : categorySignificance StartsWith "/Compromise" ①  
Node : categorySignificance EQ "/Suspicious" ②  
Node : categorySignificance EQ "/Hostile" ③
```

The following lines indicate the optimized order of conditions (indicated as nodes):

```
Optimized nodes are  
Node : categorySignificance EQ "/Suspicious" Moved from position 1  
Node : categorySignificance EQ "/Hostile" Moved from position 2  
Node : categorySignificance StartsWith "/Compromise" Moved from position 0
```

Observe how nodes were moved up and down to indicate which conditions are evaluated first. Index 0, originally profiled as the costliest, is now evaluated last.



Note: Filter conditions can have nested filters. In this case, conditions in those nested filters are also evaluated and optimized, but the node listing in the log will not distinguish in which filters the conditions are defined.

To save a filter in its optimized state:

If you are a content author, you can leverage the log information to modify your filter conditions permanently. This procedure is optional.

1. Refer to the log and re-order the filter conditions accordingly.
2. Save the filter in this optimized state. Repeat for nested filters as required.

Filters may be shared with other resources, for example, active channels and query viewers. In this case, if you manually optimize filters used by data monitors, other resources using the same filters will benefit from the change. Once filter conditions have been optimized for a data monitor and even if you don't manually change the filters, you may disable the optimization tool, described in ["Disabling the Optimization Feature" below](#).

Disabling the Optimization Feature

Optimization runs only once on event filters used by data monitors and the profiles for the evaluated conditions are stored in the log file. Each time you disable and re-enable event data monitors, evaluation and optimization runs again. When the feature is no longer needed, for example, you have modified the event filters, you can change the optimization setting to `false` in `server.properties`. While the tool is disabled, any newly-deployed data monitor will not be optimized.

Using Charts

The Console offers several chart view options for active channels and for data monitors. You can add chart views of the data in many active channels or data monitors simply by choosing a chart type from the **Format** pop-up menu in the view's lower-right corner.

ArcSight charts remain linked to the data they represent. You can immediately see a chart's events in a grid view that presents the data as charted, or filtered further using the options of the **Analyze in Channel** command. Charts use the same color for all values in a series. For example, if you are plotting successful and failed logins in a chart, successful logins as a series will have one color. Failed logins as another series will have a different color.

You can click and drag three-dimensional charts on their vertical or horizontal axes to tilt them for better viewing.



Note: Contents of charts are affected by the factors that affect active channels or data monitors, such as changing time parameters or filters. Not all charts are applicable to, or available for, all views.

Charting an Active Channel's Contents

Where: Navigator > Resources > Active Channels

1. Right-click a channel and select **Show Active Channel**.
2. In the Viewer panel, in the lower-right corner of the newly opened active channel, click the **Viewer Selector** icon to open its menu.
3. In the menu's **Chart** branch, select one of the chart types:

Chart Type	Description
Area	A horizontal chart in which bands occupy various amounts of the displayed area to indicate relevant values.
Horizontal Bar	A horizontal chart that shows changes in relative quantities, usually by time units seen as solid rectangles, over a span of time.
Line	A horizontal chart that shows changes in relative quantities, usually by time units plotted on a line, over a span of time.
Pie	A circular chart with proportional wedges for the relevant values.
Scatter Plot	A horizontal chart that shows changes in relative quantities, usually by time units plotted as separate points, over a span of time.
Stacking Area	A horizontal chart in which stacked bands occupy various amounts of the displayed area to indicate relevant values.
Stacking Bar	A horizontal chart that shows changes in relative quantities, usually by time units seen as stacked solid rectangles, over a span of time.
3D Bar	A corner-anchored graph with height, width, and depth dimensions that can show three axes of categorical and quantitative information.

4. The data in the view opens in an additional chart presentation, in the chosen format, within the active channel.
5. Click the **Layout** icon in the channel's lower-right corner to change the visual arrangement (tabbed or tiled) of the views within the channel, if needed.

Charting a Data Monitor's Contents

Where: Navigator > Resources > Dashboards

1. Double-click a dashboard or right-click it and choose **Show Dashboard**.
2. In the Viewer panel, in the lower-right corner of an applicable data monitor, click the **Viewer Selector** icon to open its menu.
3. In the chart menu, select one of the chart types:

Chart Type	Description
Area	A horizontal chart in which bands occupy various amounts of the displayed area to indicate relevant values.
Horizontal Bar	A horizontal chart that shows changes in relative quantities, usually by time units seen as solid rectangles, over a span of time.
Line	A horizontal chart that shows changes in relative quantities, usually by time units plotted on a line, over a span of time.
Pie	A circular chart with proportional wedges for the relevant values.
Scatter Plot	A horizontal chart that shows changes in relative quantities, usually by time units plotted as separate points, over a span of time.
Stacking Area	A horizontal chart in which stacked bands occupy various amounts of the displayed area to indicate relevant values.
Stacking Bar	A horizontal chart that shows changes in relative quantities, usually by time units seen as stacked solid rectangles, over a span of time.
3D Bar	A corner-anchored graph with height, width, and depth dimensions that can show three axes of categorical and quantitative information.

4. The data in the monitor switches to a chart presentation.

For data monitors, the **Chart Showing Priorities** submenu offers many of these same charting options, but with graphic elements such as pie wedges or bar segments that distinguish their priority-level components.

For more about the format tools available for dashboards () , see ["Monitoring Dashboards" on page 252](#).

For more about custom view dashboards, see ["Using Custom View Dashboards" on page 261](#).

For more about working with dashboards, see ["Using Dashboards" on page 252](#).

Exploring the Events Behind a Chart

To see a grid view of the events behind an active channel's chart, double-click the section of the graphic that represents those events. To filter those events further, right-click the relevant section of the chart and choose an **Analyze in Channel** command option. In charts that show

color keys, such as Events by Priority, you can also double-click a color chip to open a grid view filtered by that key.

To see an active channel grid view of the events behind a data monitor's chart, double-click the section of the graphic that represents those events, or right-click and choose **Show Details**, or choose **Show Detailed Channels** to see a view for each of the chart's components.

Using Query Viewers

Query viewers are a type of resource used for defining and running SQL queries. You can run query viewer results directly or add query viewers to a dashboard. For more information about defining query viewers and adding them to the dashboard, see ["Query Viewers" on page 340](#).

Graphing Attacks

You use graphic analytics to quickly identify high-volume attackers or targets at a glance. You can locate and typify cascading attacks (for example, worms and viruses), and isolate and analyze events involving interactions between two or more devices (for example, threat discovery).


- The event data you visualize can be **static** (a snapshot of the selected events) or **live** (continuously updated with specified real-time event data). You create static graphs by selecting certain event data out of a source and displaying it as a graphic. See ["Creating Static Event Graphs" below](#).
- You create live graphs using a graphic data monitor type or an active list. See ["Creating Live Event Graphs" on the next page](#).
- See ["Event Graph Notes" on page 289](#) for descriptions of the graphical elements on an event graph.
- See ["Configuring Event Graphs" on page 57](#) to set or change your event graph preferences.

Creating Static Event Graphs

Where: Event grid, data monitor, or Event Inspector panel

1. Select an array of events in a grid, data monitor, or event inspector.
2. Right-click the selected set and select **Event Graph** or **Geographic View**.

The Viewer panel displays the selected events in a new view, using the selected style.



Creating Live Event Graphs

You can create live event graphs using either data monitors or active lists.

Using a Data Monitor

Where: Navigator > Resources > Dashboard > Data Monitors > *event graph* or *geographic event graph* data monitor

- Right-click the data monitor and select **Add to Dashboard As>Geographic Graph** or **Graph**.
- Alternatively:
 - a. Right-click your personal Data Monitors folder in the Navigator and select **New Data Monitor**. A Data Monitor Editor panel opens.
 - b. In the **Data Monitor Type** drop-down list, select **Event Graph** or **Geographic Event Graph**.
 - c. Define the graphic data monitor as described in ["Creating a Data Monitor" on page 266](#).

The Data Monitor Editor has certain attributes for these types:

Attribute	Usage
Max Event Count	The number of most-recent events to show. Events older than this are discarded.
Event Node Identifier	The fields that are available to use to uniquely identify the event type in a transaction.
Availability Interval	The number of seconds for the interval between updates to the graphic.

Attribute	Usage
Show Source-Target Nodes as	See "Configuring Event Graphs" on page 57.
Source Node Identifier	See "Configuring Event Graphs" on page 57.
Target Node Identifier	See "Configuring Event Graphs" on page 57.
Show Event Nodes	See "Configuring Event Graphs" on page 57.

For geographic event graphs:

Geographic Event Attributes

Attribute	Usage
Max Event Count	The number of most-recent events to show. Events older than this are discarded.
Availability Interval	The number of seconds for the interval between updates to the graphic.


Using an Active List

Where: Navigator > Resources > Lists > Active Lists

1. Right-click the appropriate active list, and select **Show Entries**.
2. On the Details tab, select one or more events.
3. Right-click the selected events, and select **Event Graph**.
4. Configure the Event Graph window settings. For information on each of the settings, see ["Configuring Event Graphs" on page 57.](#)
5. Click **OK**.

Event Graph Notes

Link-analysis visualizations are chart-like or logically oriented. Geo-spatial visualizations are map-based or physically oriented. Node size indicates increasing event volume.



Each event is composed of the event node itself (a turquoise circle) and its connected source node (red square) and target node (white square) device assets. The source and the target may be the same asset.

Blue squares indicate a combined source and target node (a “point event”). Pink nodes indicate IP addresses that are worm or virus infection sources for other nodes.

Point events occur on a single host; for example, a syslog entry for a running process. They graph as IP address nodes that loop to an event node and back.

In geo-spatial displays, source and target location plotting is based on the physical addresses registered for IP addresses. ArcSight includes standard plotting information for this purpose. The addresses are plotted against a world map that you can zoom in or out. All the specific location data that supports this feature also appears as attributes in the Event Inspector.

You can modify the way graphs plot events, choosing to keep the source-event-target visual relationships compact, or to emphasize unique sources, targets, or both in order to more easily clarify the nature of attacks or situations.

Chapter 9: Selecting and Investigating Events in Active Channels

An active channel is a grid view in which there is a row for each event. In an active channel you can select which events you want to investigate. After selecting one or more events in the channel, you can perform several analysis and authoring tasks.

Selecting Events in the Active Channel

To select fields to investigate:

Click an event or **Ctrl+click** a set of events. To select a range of events, click one event and **Shift+click** the event at the end of the range.

To invert your field selection:

Select one or more events in the channel, right-click and select **Invert selection**.

To select events with matching cells:

Select a cell in an event, right-click and select **Select events with matching cell** to see if other events in the channel have matching cell values.

Showing Event Details and Rule Chains

Rule-based correlation events are those generated by a triggered ArcSight rule as a reaction to an original sensor-generated event. In other words, an event concerning an event. You recognize correlation events in active channels by their red **Flash** icon . To mask active channels so they show **only** correlation events, select the check box at the top of the channel's left-most column.

To display event details:

1. In an active channel, select an event.
2. Right-click and select **Show event details**. The event's details appear in the Event Inspector.

More:

- Some system operations, for example, audit event generations, are done on behalf of a special system user called 1ROOTUSER. When you are investigating event details, you might see a user ID with this value. This user ID is valid and intended for internal use only.
- When you apply an actor field set to an event being displayed in the event inspector, you may experience an extended load time.

To display simple event rule chains:

1. In an active channel, select a correlation event. You recognize correlation events in active channels by their red **Flash** icon .
2. Right-click and select **Rule options > Simple chain**.

To display detailed event rule chains:

1. In an active channel, select a correlation event.
2. Right-click and choose **Correlation options > Detailed chain**.

The events leading up to the correlation event appear in the Description panel at the top of the Inspector.

3. Click any event in the chain to see its details on the panel, below the description.

To view forwarded correlation events and their correlated (base) events:

Based on your requirements, you can configure the Forwarding Connector to send only the correlation events without the correlated (also referred to as base) events from one source Manager to a destination Manager. You can also configure the connector to include the correlated events automatically, whenever a correlation event is forwarded.

After the correlation and correlated events are forwarded, you can view them on the destination through the Event Inspector. Make sure the source Manager is configured correctly according to the instructions in the *Forwarding Connector Configuration Guide*.

To display correlation-event rules:

1. In an active channel, select a correlation event. You recognize correlation events in active channels by their red **Flash** icon .
2. Right-click and select **Correlation options**, then **Show triggering resource**.

The rule or resource that triggered the correlation event is selected in the Navigator panel's Rules resource tree and that rule appears in the Rules Editor.

To execute or clear rule actions:

1. In an active channel, select a correlation event. You recognize correlation events in active channels by their red **Flash** icon .
2. Right-click and select **Rule options**, then **Clear Rule Actions** to clear all actions associated with this rule.

For more information, see ["Managing Rule Actions" on page 532](#).

To launch event details in a browser:

1. In an active channel, right-click an event and choose **Show event details**.
2. In the condition table of the Event Inspector, right-click and choose **Launch Event Details in Browser**.

A Web browser opens with the selected event's details.

To hide empty rows in the Event Inspector:

1. In an active channel, right-click an event and choose **Show event details**.
2. In the condition table of the Event Inspector, right-click and choose **Hide Empty Rows**.

Running Recon Searches

You can conduct event searches from ESM using Recon. Refer to the *Recon User's Guide* for details.

Integration with Recon requires specific browser versions. See the *ESM Support Matrix* for details.



Tip: Events that are not within the latest half hour time range are not displayed in Recon as search results. Events must be within the time range shown in the ArcSight Console in order to be displayed.

Prerequisite:

ESM must be integrated with a deployment of Recon and ArcSight Transformation Hub. If the required setups are done correctly, the following integration commands are enabled on an active channel and from the event details on the Inspect/Edit panel:

- Recon
- Recon (Multiple Fields)

**Note:**

- Not all ESM fields are supported in Recon searches. These unsupported fields will appear disabled for selection.

To search a single field:

1. Open an event viewer such as an active channel, or view an event's details in the Inspect/Edit panel.
2. Right-click a row and select **Recon**.

Recon displays the search results in your preferred browser.

To search multiple fields:

1. Open an event viewer such as an active channel, or view an event's details in the Inspect/Edit panel.
2. Right-click a row and select **Recon (Multiple Fields)**.

The Recon panel opens, displaying a list of fields supported for the search. The list is based on the columns that are available on the channel.



Tip: If you know the field name and prefer not to scroll through the list to locate it, enter the name in **Search Fields**. Enter the first few characters until the matching field is selected.

3. Click **Add** to add the field to the Selected Fields pane. Select up to five fields.
4. Click **OK** to begin the search in Recon.

Recon displays the search results in your preferred browser.

See also:

- The topic, ["Using the Recon Integration Commands" on page 686](#).
- *Recon User's Guide*

Investigating Session Events

Session events are captured in session lists, described in:

- ["Identity Correlation" on page 565](#)
- ["Managing Session Lists" on page 500](#)

This topic explains how to use the 's **Analyze in Channel >Session Events** command to refine and explore channels contextually, using attributes of the events already being displayed in active channels.

You can investigate session list entries two ways:

- Filter the set of entries based on the attributes of a particular entry, or
- Create an Investigation Channel that contains only the entries that match one or more attributes of the initial session list entry.

To investigate a session event:

Where: Navigator > Resources > Lists > Session Lists tab

1. Right-click a session list and select **Show Entries**.
2. In the Viewer panel, click an entry that bears investigation.
3. Right-click the selected entry. The menu includes commands to **Create Channel** and **Add Condition to Channel Editor**. The details of each command will vary based on which column you right-click.

For example, if you right-click a Source IP column containing the value 192.0.2.0, the choices are:

- Create Channel (Source IP = 192.0.2.0)
- Create Channel (Source IP != 192.0.2.0)
- Create Channel >
- Add Condition to Channel Editor (Source IP = 192.0.2.0)
- Add Condition to Channel Editor (Source IP != 192.0.2.0)
- Add Condition to Channel Editor >

The sub-menus (indicated by the >) offer similar choices for all the other columns of the Session List entry.

If you select **Create Channel**, a channel is added to the Viewer panel. If you select **Add Condition to Channel Editor**, a channel editor will open in the Inspect/Edit panel.

For more information about creating and using views for investigation, see ["Investigating Views" on page 233](#).

Collaborating on Events (Event Annotation)

You can use workflow-style annotation to collaborate with other users in analyzing or reviewing selected events.

When you are annotating, you can make collaboration-stage changes to just the event you originally selected, or have that change also affect a larger set of similar events that should also be carried forward in the review process.

The central tasks in annotating events for collaborative analysis are assigning them to yourself or another user, then assigning them to one of the available sequential workflow stages (dispositions). While ArcSight comes with a default set of stages, your enterprise will very likely have customized these stages and created new ones.

Compare collaborative annotation to cases, which are a more formal way to track sets of events that are under investigation.

Related topics:

- ["Annotating an Event" below](#)
- ["Viewing Annotations for an Event" on page 299](#)
- ["Creating or Editing Stages" on page 300](#)
- ["Case Management and Queries" on page 624](#)

Annotating an Event

Purpose: To flag one or more events that enables tracking those events through a workflow. See the topic, "Annotations" in *ESM 101*.

Where: Viewer panel displaying an active channel of events

Procedure:

1. Select one or more events in any active channel. If not already annotated, you can start a collaboration cycle.
2. Right-click the events and select **Annotate Events** or press **Ctrl+T**.
3. In the Annotate Events dialog popup, set or change the events' Annotations fields, as

described below.

Event Annotation Fields

Field	Usage
Stage	Click this field to choose a different disposition state for the events' collaboration cycle. The default stage is [Queued] and available stages run from Initial to Closed. If you created your own stages as described in "Creating or Editing Stages" on page 300, these custom stages would be displayed here. Setting the event's Stage through a rule action: You can also automate the setting of the selected event's stage through the <code>Set Event Field</code> rule action. Every time the rule triggers, the stage set by the rule will take precedence over the stage setting done by this manual event annotation. Other events marked as similar to this event are affected in the same way: their stages will be set by the <code>Set Event Field</code> rule action. If you want to override the rule action, add this statement to the <code>server.properties</code> file: <pre>mark-as-similar.override-annotation-stage=true</pre> For the instructions on how to edit the <code>server.properties</code> file and implement your change, refer to the topic "Managing and Changing Properties File Settings" in the ESM Administrator's Guide. See also Set Event Field in the Rule Actions Reference topic.
Assign to	Click this field to choose an ESM user to take the next step.
Is Reviewed	This read-only field tells you whether this event has been reviewed.
In Case	This read-only field tells you whether these events are already part of an ESM case. If they are, you have more ways to track their disposition. See "Viewing a Case's Events in a Channel" on page 643 for related information.
Correlated	This read-only field tells you whether these events are part of a correlated event chain. If so, you can learn more through the rules authored to control that chain of correlation. Note: You can configure the ArcSight Forwarding Connector to send correlation events along with the correlated base events from a source Manager to a destination Manager. However, the forwarded base events display the Correlated flag only in memory, which you can view on an active channel. If you want the forwarded base event's Correlated flag to persist in the database, set this property in <code>server.properties</code> on the source Manager: <pre>logger.base-event-annotation.enabled = true</pre> For instructions on how to edit the <code>server.properties</code> file, refer to "Managing and Changing Properties File Settings" in the ESM Administrator's Guide.
Hidden	This read-only field tells you whether these events are hidden from all but the assigned users of this stage.
Closed	This read-only field tells you whether the investigation of these events has been marked as closed. Closed events may no longer be visible to interested parties through active channels, for example.

4. Add information in the **Comments** field as needed to clarify the collaborative process.

5. To have your changes also affect related events, use the **Mark Similar Events** fields, as described in [Mark Similar Events Fields](#).
6. Click **OK** to update the event.

Mark Similar Events Fields

This topic is continued from ["Annotating an Event" on page 296](#).

Event “similarity,” for collaboration purposes, is defined as a combination of time constraints and having certain key event attributes in common. For example, you could apply a collaboration change to additional events received in the future on the basis of those events having the same Attacker value and having occurred within the last two days.


After the time constraints have passed, the common event attributes return to default settings for events marked as similar to a specified event.

Purpose: To specify one event attribute that will be the basis for event similarity when annotating fields.

Where: Annotate Events popup dialog as a continuation from the procedure, ["Annotating an Event" on page 296](#).

Procedure:

1. At the bottom of the Annotate Events dialog, click the plus sign (+) next to **Mark Similar Events** to expose the available fields.



2. Set the other attributes as shown:

Similarity Field	Usage
Time Constraints	Choose a bracketing combination of Start Time and End Time or Duration to determine the scope of the constraints. Note: After this time constraint has passed, the events' stage annotation reverts to [Queued].
Start Time	Date and time values to set the beginning of a time-constraint window. Choose from the drop-down menu of expressions or click the ellipsis button to set exact times.
End Time	Date and time values to set the end of a time-constraint window. Choose from the drop-down menu of expressions or click the ellipsis button to set exact times.
Duration	The length of the time window, relative to a Start Time or End Time, when using Duration as a time constraint.
Criteria	A menu of key event-attribute characteristics you can use to define similarity. The text box below specifies the criteria being set. For example, do you want similarity based on the same event name,? Each selection displays the corresponding description. For example, if the original event you selected was called Monitor Event and you selected Same Name as criteria, then that event name is reflected in the text box below.

3. Click **OK** to save your entries and to complete your event annotation configuration.

From now on, any event matching the criteria will be assigned to the same Stage attribute as specified in ["Annotating an Event" on page 296](#). Pay attention to the description about setting the Stage and how it can be overridden by a rule action.

Annotation Preservation

With the CORR-Engine, when the day's events are archived at the end of the day, the archive is permanent. If, while the archive is still online, you make changes or additions to event annotations, they are preserved as a supplemental archive when the archive goes off line at the end of the retention period. If you reactivate an offline archive and make more annotation changes, they are only preserved until you deactivate the archive, at which time these annotation changes are deleted. Refer to the [ArcSight Command Center Users Guide](#) for information on retention periods and archives.

Viewing Annotations for an Event

Annotations on an event are displayed in the **Annotations** tab of the Event Inspector when that event is selected.

To view the annotations for an event:

1. Right-click an event in an active channel (such as an active channel or active list) and select **Show Event Details** to display the Event Inspector.
2. In the Event Inspector, click the **Annotations** tab.

The tab displays the event's timestamp, the user name associated with the event, the stage, and flags. See ["Annotating an Event" on page 296](#) for related information.

Creating or Editing Stages

For a description, see the topic "Stages" in *ESM 101*.



Caution: Keep stages provided as standard content in the given folders and do not move them into another folder. Standard content stages are Closed, Final, Flagged as Similar, Follow-up, Initial, Monitoring, Queued, and Rule Created.

Where: Navigator > Resources > Stages

1. If you are creating a stage, right-click the **All Stages** group and select **New Stage**.
If you are editing a stage, right-click a stage under the **All Stages** group and select **Edit Stage**.
2. In the Stage Editor, enter a name for the stage.
3. Set the fields as described in the following table:

Stage Editor Fields

Field	Usage
Subsequent stages	Select one or more stages to set as follow-on stages to this one. Events in this stage will show these other stages as options in the Stage field of the Annotate Fields dialog box.
User required	Select whether you want to prompt for a user assignment when assigning this stage. If you don't prompt for a different user, or no change is made, the current user remains in effect.
Comment required	Select whether you want to require users to add a comment when assigning this stage.
Can be skipped	Select whether this stage can be bypassed when assigning from one stage to the next.
Mark similar required	Choose whether you want events that are similar to the selected events to be automatically assigned to this stage. Similarity is scoped at assignment time through the Mark Similar Events fields of the Annotate Events dialog box you see when you choose Annotate in an active channel. Note that similarity marking applies only to subsequent events received in the future. Events already processed are not affected.

Stage Editor Fields, continued

Field	Usage
Mark similar stage	Select whether you want to use this stage as a routing mechanism for other stages in a workflow. When selected, assigning one or more events to this stage causes all following (subsequent) similar events to be automatically redirected to the chosen stage. Events already processed are not affected. Similarity is scoped at assignment time through the Mark Similar Events fields of the Annotate Events dialog box you see when you choose Annotate in an active channel. Note: With the assistance of ArcSight Professional Services, you can customize the similarity criteria selector for Mark Similar events. In this way you can have conditions that are different from the defaults. This is done with the Velocity scripting language, by modifying certain Velocity templates present on the , in the config/similarity directory. Ask your ArcSight administrator for more information or make a request of ArcSight Professional Services.
Hidden	Select whether you want events assigned to this stage to be hidden from all but the assigned users (True), left visible to everyone (False), or to leave the current visibility unchanged (Ignore).
Closed	Select whether you want events assigned to this stage to be marked as closed to investigation (True), not marked as closed (False), or left in their previous state (Ignore).

4. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
5. Click **Apply** to save your changes and keep the editor open, or click **OK** to save and close.

Working with Event Payloads

An event payload is the information carried in the body of the event's network packet, as distinct from the packet's header data. From the , you can search, retrieve, view, save to a file, or discard event payloads.

The first step in handling event payloads is to be able to locate payload-bearing events among the general flow of events in an active channel.

Events payloads are identified by the Payload ID icon:

**To find payloads:**

1. In an active channel, right-click a column header and choose **Add Column>Device>Payload ID**.
2. Look for events showing a Payload ID in that column.

To retrieve payloads:

1. In an active channel, double-click an event with an associated payload.
2. In the Event Inspector, click the **Payload** tab.
3. Click **Retrieve Payload**.

To preserve payloads:

You can select to preserve the payload for an event in either of two ways:

- In an active channel, right-click an event with an associated payload, choose **Payload**, then **Preserve**.
- Or**
- In the Event Inspector, click the **Payload** tab, then **Preserve Payload**.

To discard payloads:

In an active channel, right-click an event with an associated payload and choose **Payload**, then **Discard Preserved**.

You can also use the Event Inspector, as follows:

1. In an active channel, double-click an event with an associated payload.
2. In the Event Inspector, click the **Payload** tab.
3. Click **Discard Preserved Payload**.

To save payloads to files:

1. In an active channel, double-click an event with an associated payload.
2. In the Event Inspector, click the **Payload** tab.
3. Click **Save Payload**.
4. In the Save dialog box, navigate to a directory and enter a name in the **File name** text field.
5. Click **Save**.

To view payloads in other viewers:

1. In an active channel, double-click an event with an associated payload.
2. In the Event Inspector, click the **Payload** tab.
3. Click **Launch External Payload Viewer**.
4. View the payload using the **Preferred Payload Viewer** and **Text to PCAP Converter**, specified in the 's **Edit>Preferences>Programs** panel.

Exporting Data Fields to a CSV File

You can export a set of data fields into a comma-separated values (CSV) file.

To export to a CSV file:

1. In the channel, select one or more events.
2. Right-click and select **Export > Events in Channel**.
3. On the Export Events file browser, navigate to the location where you want to save the CSV file, then enter or select options for these fields:

File Name	Enter a file name for the CSV file. Note: The file name extension is not required; the csv extension is added automatically when the file is created.
Files of Type	Select Comma separated values (*.csv).
Export Data Options	For Rows, you have two options: • If you select All in channel, all events in the channel will be exported to the CSV file. • If you select Selected rows only, only those rows highlighted for the right-click operation will be exported to the CSV file. • The default for Columns is the Export field set. Keep the default or select other field sets from a list of All Field Sets. • For Destination, select Local CSV File.

4. Click **OK**.

To limit the export to fields visible in the channel:

Purpose: The default “Export” field includes a large number of columns. Unless you have a pressing need to export all these fields for channel events, you might want to modify the export. Exporting a large field set for a large event set could be time- and resource-consuming.

- If the channel is unmodified from its default (you have not added or removed fields), you can select the channel’s default field set on the file export option. To find the default field set name, edit the channel and look at “Default Field Set” name or right-click any column header in the channel and choose **Field Set > Selected Field Set**. The default field set will be selected. (For example, for /All Active Channels/ArcSight System/Core/Live active channel, the default field set is Standard-MgrRcpt. Selecting this field set on the export will give you that set of columns in the CSV file.)
- If you have modified the channel from its default (added or removed fields), you can save it as a custom field set and then choose your custom field set on the export dialog. To save a

custom field set, right-click anywhere on the column headers in the active channel and choose **Field Sets > Save As**. On the Field Sets Selector, navigate to the group you want, name the new field set and click **OK**. Now it will be available to choose from on the export dialog.

The Export field set itself is also customizable. If you are sure you always want exported events to include a limited set of fields, you can edit the Export field set. See ["Creating a Field Set" on page 601](#) and ["Editing a Field Set" on page 607](#).

Exporting Events to ServiceNow[®]

You can export ESM events and create them as tickets in ServiceNow[®]. When you export an event, ESM creates an audit event with the following attributes:

- **Name** - Created Ticket on External System
- **device event classid** - externalticket:100
- **device event category** - /externalticket/snow
- **device custom string 1** - contains all exported event IDs, concatenated into one string
- **device custom string 2.resource URI** - specifies the path of the rule that triggered the export
- **File ID** - contains a hyperlink to open the ticket in ServiceNow[®]

ServiceNow[®] manages the session and session expiration between the ArcSight Console and ServiceNow[®]. If your ITSM session expires and you are not done with your tasks, log in again.

Ensure that your environment meets the following prerequisites before you attempt to export an event:

- Verify that ESM is configured to integrate with ServiceNow[®]. If not, run the Manager setup wizard and provide the requested information. For more information, see the [ESM Installation Guide](#).
- ServiceNow[®] must be running and you must have login privileges.

To export an ESM event:

Displaying the event editor is not required for this procedure. You can export an event whether it is locked or unlocked.

1. In an active channel, select one or more events to export.
2. Right-click the selected events and select **Export > Export to ServiceNow[®] ITSM**.
3. Use one of the following methods to log in to ServiceNow[®]:

- Use the ServiceNow® Global ID.
 - Provide your ServiceNow® credentials and click **Next**.
 - To use token authentication, select the **OAuth** checkbox and click **Next**. Enter the token and click **Next**.
4. Provide the requested information, and then click **Create Ticket**.

Getting Knowledge Base Articles

Knowledge Base articles can be associated with events, rules, or any resource. Knowledge Base articles can have links or notes to help you respond to events.

HTML-based articles are displayed on your default Web browser.

Where: Navigator > Resources > Knowledge Base > Knowledge Base Articles

To display articles from the Knowledge Base window:

Right-click an article and select **Show Article**.

To display articles from an active channel:

In an active channel, right-click an event and select **Knowledge Base**, then **Show**. Select **KB entry for cell**, **KB entry for row**, or **KB entry for column**, then the article name.

For more information about active channels, see ["Monitoring Active Channels" on page 224](#).

To display articles from an Event Inspector:

In the Event Inspector, right-click an event and choose **Knowledge Base**, then **Show Article**.

Chapter 10: Filtering Events

The Filters resource tree in the Navigator panel is pre-populated with some typical event filters you can use directly, or as templates for more specific purposes. You can create and edit your own filters and inline filters for use in active channels.

Creating or Editing a Filter

If you want to make a filter available to multiple resources, follow these instructions to define a filter resource.



Tip: Best practices for creating and using filters

- Because you can reference filters in other filters, you can create hierarchies similar to style sheets. Plan your filtering needs so you can create filters, filter groups, and filter hierarchies that will promote the most efficient and consistent analysis results.
- Learn how to use the Common Conditions Editor (CCE). Refer to ["Common Conditions Editor \(CCE\)" on page 887](#), ["Conditional Statements" on page 903](#), and ["Conditions" on page 904](#) for more information.

Where used:

- Active channels
- SmartConnectors
- Reports
- Rules
- Other filters



Note: When a filter condition is changed, the change to the condition is add to the notes on the filter, tracking the history of the condition change.

Where: Navigator > Resources > Filters

1. If you are creating a filter, right-click a filter group and select **New Filter**.
If you are editing a filter, right-click that filter and select **Edit Filter**.
2. In the Filters Editor, type in the **Name** text field.

The Sub Type is fixed to Event Filter.

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes

sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

3. Go to the **Filter** tab and define filter conditions:
 - a. In the table, scroll to a relevant event field and choose a logical operator (**Op**), enter a conditional statement (**Condition**).
 - b. Select case-sensitivity (**Aa**), and select inequality or negate (**Not**), if appropriate. Use the features described in ["Common Conditions Editor \(CCE\)" on page 887](#).
4. Repeat the above step for each condition you want to add to the filter.
 - To edit a logical operator, right-click the logical operator and choose **Edit**, then choose a logical operator and click **OK**. (For more information, see ["Logical Operators" on page 1006](#).)
 - To delete a logical operator, right-click the operator and choose **Delete**. In the confirmation dialog box, click **Yes**. The logical operator and all its condition statements are removed.
 - To delete a condition statement, right-click it and choose **Delete**. In the confirmation dialog box, click **Yes**.

**Caution:** Filter definitions (meaning the total text used in a filter's condition statements) cannot exceed 10,000 characters. If your filter uses more than 10,000 characters, create a second filter by splitting the definition, and use the **MatchesFilter** operator to combine the two.

 - To view the full conditions for the **MatchesFilter** operator, click the **Summary** tab and then click the **Expand Filter** button to display the filter conditions for debugging.
Note that in this case, the display of the **MatchesFilter** full logic does not display the sub-filter of the matched filter. Full logic is displayed only for the first level of matched filter conditions.
5. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
6. Click **Apply** below the Inspect/Edit panel to update the filter or click **OK** to add the filter to the resource tree.

Creating and Editing an Inline Filter

Purpose: To filter events for an active channel. This inline filter will only be for the exclusive use of that active channel.



Tip: Steps to create an inline filter are summarized here. For more details and examples, see also ["Filtering Active Channels with Inline Filters" on page 234](#).

In any active channel grid view you can use the fields of the grid's top line to select filtering event-attribute values for certain columns, which will be used with implied AND operators to apply temporary filters and use the grid's bottom line to select filtering event-attributes values which will use OR operators.

These filters are **not** retained with the active channel, but you can give the revised channel a name and save it through the Active Channel Editor.



Note: You cannot select a grayed-out column to include in your filter. Grayed-out columns have either variables or they are a custom column.

Where: Resources > Navigator > Active Channels

1. In the Active Channels resource tree, open the channel to which you want to add an inline filter.
2. In the Viewer panel, go to **Inline Filter** and click **No Filter**. This opens the inline filter pane.
3. Select the parameters based on the active channel's columns. For example, if you want to filter through the column called Name, enter a name value. Click **Apply**.
4. To highlight all matching events for your filter, select the **Highlight** check box. Highlighting allows you to preview the events that match your filter prior to saving the filter. Click **Apply** to activate the inline filter.

You can specify the highlight color by clicking the drop-down picker and select your color.

5. To add or delete rows to the inline filter table, click + (plus) or click - (minus).

To create and manage multiple inline filters, click the + button next to the Highlight options under the inline filters to add filter definition rows. (Click the - button to remove filter rows.) The potential uses of multiple inline filters are extensive, but essentially this provides a means of creating a filter with complex conditions, inline in an active channel. For example, in the Name column for an event, you could specify that the event name contains `ActiveList` on the first filter row and that the name does not contain `Successful`. You could extend this filter by specifying what you are looking for in some of the other fields or even add more qualifiers on the Name field. All fields can be narrowed down in this way, using multiple filter definition rows.

Applying Filters

This topic discusses how to use filters in these resources:

- Active channels
- SmartConnectors

- Reports
- Rules
- Other filters as part of the filter hierarchy

To add filters to resources:

You apply existing filters to other resources by referencing them in those resource editors.

1. Right-click a resource in the Navigator panel such as a filter or rule and choose **Edit <resource>**.
2. Click the editor's **Conditions** tab if it isn't already at the front.
3. In the Inspect/Edit panel, click the **Filters** button and select a filter in the Filter Selector dialog box. The selected filter becomes a new condition line in this resource's filter.
4. Click **OK** or **Apply** to save the resource's definition including its new filter reference.



Note: You can use hierarchies of filter references (including filters within filters) to better manage them, similar to style sheets.

To apply resources as filters to active channels:

You can quickly apply or test the effects of using particular SmartConnectors, assets, categories, zones, vulnerabilities, customers, stages, or filter resources as conditions to filter active channels. These filters make the referenced resource a condition for the channel in use. You can choose to make the condition exclusive or additive.

1. Open the channel to filter in the Viewer panel.
2. In an applicable resource tree in the Navigator panel, right-click an item and choose **Set as current filter** or **Add to current filter**. The filter change takes effect automatically and the channel's header immediately shows the new filter condition exclusively (set as) or as an addition (add to).

Also, you can use drag and drop to apply a filter to an active channel. This action is the same as choosing **Set as current filter** or **Add to current filter**.

3. You can click the filter description in the channel's header to open the filter in the Active Channel Editor.

To remove a filter condition from a resource:

You use the Filters tab of a resource's editor to change or remove any filters that affect it.

1. In the Navigator panel, right-click the filtered resource and choose **Edit <resource>**.
2. In the Inspect/Edit panel, click the **Filter** tab of the resource's editor.

3. In the Conditions editor, right-click the statement that imposes the condition you want to remove and choose **Delete**.
4. Confirm the deletion and click **Apply** to restart the channel.



Note: Using enforced filters

You can designate the events which users can see by adding filters to the user group's Events tab on the ACL Editor. Refer to ["Adding or Removing Enforced Filters" on page 113](#).

Moving or Copying Filters

1. In the Filters resource tree, navigate to a filter and drag and drop it into another group.
2. Choose **Move** to move the filter, **Copy** to make a separate copy of the filter, or **Link** to create a copy of the filter that is linked to the original filter.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

Deleting Filters

Be careful about deleting filters. Verify first if these filters are being used by other resources. See ["Applying Filters" on page 308](#) for more information.

To delete a filter resource:

1. In the Filters resource tree, right-click a filter and choose **Delete filter**.
2. In the dialog box, click **Yes**.

To delete an inline filter:

In the active channel header, right-click the inline filter definition and choose **Remove Inline Filter**.

Or:


1. Click the Inline Filter edit button () to display the inline filter parameters.
2. Click inside the inline filter grid.

3. Click **Clear** then click **Apply**.

The active channel is restored to its original unfiltered view.

Debugging Filters to Match Events

You can use a filter debugger to test whether a selected filter matches a certain type of event and, if there are mismatches, to determine which filter conditions are not matching the event details.



The filter debugger compares the conditions in a selected filter with the metadata that describes the selected event to determine whether the filter would capture such events. The filter definition is displayed to show the results of this comparison.


- If the selected filter matches the event, the filter definition shows no errors or mismatches.
- If the filter does not match the event, the filter definition highlights the mismatches between the filter conditions and the selected event with red-highlighted **Xs**.



Note: Red highlighted **Xs** in a filter as a result of filter debugging on an event *do not necessarily indicate* that the filter is *invalid*. Red highlights are shown here only to indicate where the selected filter does not match the selected event.

To debug a filter against an event:

1. Select an event in the viewer in an active channel against which you want to test a filter.
2. Right-click and select **Debug Filter**.




3. In Filter Selector, select the filter you want to test.

In a few moments, the Debug Filter dialog displays the filter's event conditions with applicable indicators, as follows:

- If the selected filter matches the event, the Debug Filter dialog displays the selected event with green checkmarks.
- If the filter does not match the event, the Debug Filter dialog displays the selected event with red Xs.

The following example shows the debug filter results that found a combination of matches and mismatches:



In the example, you see an OR condition comparing two events. The evaluation process found no match in the first set of conditions and found matches in the second set of conditions. The particular filter we selected, ArcSight Internal Events, happens to have a third condition to match another filter, ASM Events. However, this third condition was skipped since evaluation stopped once matches were found.

For more information about using the Event Inspector to investigate events, see ["Inspecting and Editing" on page 68](#) and ["Event Inspector" on page 994](#).

See also ["Creating or Editing a Filter" on page 306](#) and ["Applying Filters" on page 308](#).

Importing and Exporting filters

To import and export filters, use the packages feature as described in ["Managing Packages" on page 730](#). Portable ArcSight packages can automatically manage dependencies across resources and other packages.

For information on how to import and export filters on SmartConnectors, see ["Importing and Exporting SmartConnector Configurations" on page 205](#) (especially the topics on ["SmartConnector Filters" on page 206](#) and ["Adding SmartConnector Filter Conditions" on page 192](#)).

Managing Filter Groups

Filter groups are created to store similar groups or filters in a single location. Groups can be created within groups to meet enterprise needs. When a group is created within a group, the new group inherits the existing group's access control list (ACL).

 **Caution:** Do not exceed more than 10,000 resources in a group.

Groups and filters can be managed with drag and drop functionality. You can move or copy groups and filters into other groups. If a group is deleted, the filters within that group are also deleted.

 **Note:** To copy multiple resources at once, use Copy and Paste. You can drag and drop only one resource at a time.

Where: Navigator > Resources > Filters

To create filter groups:

1. Right-click a group and choose **New Group**.
2. In the Name text field, type in a name.
3. Press **Enter**.

To rename filter groups:

1. Right-click a group and choose **Edit Group**.
2. In the Name text field, rename the group.
3. Press **Enter** and click **OK**.

To edit filter groups:

1. Right-click a group and choose **Edit Group**.
2. In the Group Editor, edit the **Name** and **Description** text fields, and press **Enter** after each.
3. Click **OK**.

To move or copy filter groups:

1. Select a group and drag and drop it into another group.
2. Select **Move** to move the group, **Copy** to make a separate copy of the group, or **Link** to create a copy of the group that is linked to the original group.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete filter groups:

1. Right-click a group and choose **Delete Group**.
2. In the dialog box, click **Yes**.

Investigating Views

This topic explains how to use the Console's **Analyze in Channel** command to refine and explore channels contextually, using attributes of the events already being displayed in grid views.

The **Analyze in Channel** command uses these attributes, and the values found in their events, to automatically formulate simple filters or conditions.

When you create or refine a filter through **Analyze in Channel**, the Viewer panel automatically opens a new view of the channel with the filter applied. You explore the filter's effect in this view. You can keep the view by saving the channel under a new name, or discarding it by right-clicking in the grid and choosing **Close**.

When you use **Analyze in Channel** to add a condition to a resource editor such as Rules or Filters, the condition appears in the editor panel where you can modify it or click **Apply** to put it into effect.

The new or modified views you generate with the **Analyze in Channel** command can be grids, or you can choose to display them in applicable chart formats using the **Viewer Selector** icon in the lower-right corner of the Viewer panel.

To learn more about the event attributes these options use, see ["Data Fields" on page 910](#).

Using an Event Attribute to Show a New Filtered View

These options completely control the new view created, ignoring the filter in the original view. You most often use them to test and explore.

In a grid view, right-click an attribute (column) in an event listing and choose **Analyze in Channel**, followed by one of these options:

Option	Use
Create Filter [Attribute=Value]	Show only those events in which the selected attribute matches the value in the selected event.
Create Filter [Attribute!=Value]	Show only those events in which the selected attribute does not match the value in the selected event.
Create Filter [List of Related Attributes=Value, !=Value]	When the selected attribute is of a type that has related attributes, choose to show only those events that do (or do not) match one of the related attributes on the additional menu. Generally, attributes are considered related if they share a common focus such as IP addresses.

Refining a Filter with an Event Attribute

These options open a new view that uses a version of the prior filter modified to include the new filter component just selected. You usually apply these as part of a filter-refinement process.

In a grid view, right-click an attribute (column) in an event listing and choose **Analyze in Channel**, followed by one of these options:

Option	Use
Add [Attribute=Value] to Filter	Show only those events that match both the prior and new filter elements.
Add [Attribute!=Value] to Filter	Show only those events that do not match both the prior and new filter elements.
Add to Filter [List of Related Attributes=Value, !=Value]	When the selected attribute is of a type that has related attributes, choose to show only those events that do (or do not) match one of the related attributes on the additional menu. This filtering element is applied in addition to any other already present. Generally, attributes are considered related if they share a common focus such as IP addresses.


Filtering Out ArcSight Events

You can modify existing filters to refine your view to show only the events you want to see. Suppose you have an active channel that includes both system events and non-system events, but you want to see only the non-system events.

To modify the filter of the channel:

1. Double-click the filter in the channel header to get the channel editor.
2. Click the **Filter** tab in the channel editor.
3. Add this condition to the filter (with an AND):

```
!=NOT MatchesFilter("/All Filters/ArcSight System/Event Types/ArcSight Internal Events")
```



To create or customize active channels in other ways, follow this same approach. Find a filter that does what you want and add condition statements to filters for a channel. Or, as in the example above, find a filter that does the opposite of what you want, add it to a channel, and negate the condition statement as shown above. Since we wanted to limit the channel to show only non-ArcSight events, we found the ArcSight Events filter, added the ArcSight Events condition to a channel, and negated it to get the effect of filtering out all ArcSight events.

Adding an Event Attribute to a Filtering Condition

The **Add condition to editor** options apply to the editor in the Inspect/Edit panel that currently has focus. If no editor is open, the default target is the Filters Editor.

To add an event attribute:

In a grid view, right-click an attribute (column) in an event listing and choose **Analyze in Channel**, followed by one of these options:

Option	Use
Add Condition [Attribute=Value] to Editor	In the current editor, insert a new condition in which the selected attribute matches the value in the selected event.
Add Condition [Attribute!=Value] to Editor	In the current editor, insert a new condition in which the selected attribute does not match the value in the selected event.
Add Condition to Editor [List of Related Attributes=Value, !=Value]	When the selected attribute is of a type that has related attributes, add a condition to the current editor using the available list of attribute-value pairs that do (or do not) equate. Generally, attributes are considered related if they share a common focus such as IP addresses.

To remove a condition from the editor:

Right-click it and choose **Delete**.

When you are using these options to affect a view that is subject to the editor in use, click **Apply** or **OK** in the editor to put the condition into effect.

Contextual filters (in contrast to conditions) are temporary unless you save the modified view as a named active channel. Condition statements are saved with their relevant editors.

Modifying Views

This topic covers the use of inline (filtering statements are in the grid itself) grid view filtering options. The inline filter is the row of blank event values you see at the top of any grid in the Viewer panel.

Inline filtering directly affects the current view. Changes you make to a grid view by inline filtering also apply to any other versions of the view you open such as its applicable chart types. Inline filters are temporary unless you save the modified view as part of a named active channel.

To modify a view with inline filters:

1. Clicking the inline fields at the top of view columns.
2. Select an event-attribute value to use as a constraint. When you select multiple fields, they automatically form AND conditions.
3. Click the **Checkmark** icon to apply your filter selections.

To undo an inline filter:

1. Click any of the filter fields in the top line of the grid view to show the inline filter control buttons.
2. Click the **X** (clear) button to remove the current filter elements and restart the view.

For details on working with filters and inline filters, see ["Creating or Editing a Filter" on page 306](#) and ["Filtering Active Channels with Inline Filters" on page 234](#).



Tip: If you want the modified view as a permanent view, use the Navigator panel's Active Channel resource tree to open the view's channel in the Active Channel Editor. Then save the channel with a new name.

Chapter 11: Queries


A query is an ArcSight resource that defines the parameters of the data you want to report on derived from an ArcSight data source. The result of the query then becomes the basis for one or more ArcSight report and trend. The Query Editor is a component of ArcSight Reporting resource tools.

Queries are used in:

- **Reports.** See ["Building Reports" on page 388](#) for more information.
- **Query viewers.** Additionally, if you want to run quick SQL queries for monitoring and analysis outside of the reporting resource, you can use query viewers. You can add query viewers to dashboards and generate simple reports on query viewer results. For information on query viewers, see ["Query Viewers" on page 340](#).
- **Trends.** See ["Building Trends" on page 467](#).

How Queries Work

As a data source, queries can use the database of events, actors, modeled network objects (assets), cases, notifications, case-sensitive session lists or active lists, or data gathered from a trend.



In a query, you select the data fields you want to report on, specify any additional functions you want run on them (such as sum, average, and so on), and any sort or group-by conditions you want to add, such as grouping results by source address, zone, or priority.

Using Queries and Trends Together for Reports

A query can be used as the primary data source for a report. Or, a trend (based on one query) can be used as the data source to another query that further refines the initial query result. A collection of trend queries (queries that use trends as their data source) can provide focused views of a data set which can then be fed into a single report or multiple reports.

For a more detailed description of the relationships you can build between queries and trends for reporting, see the ["Query-Trend Relationships in Reporting" on page 468](#).



Tip: When running reports or trends:

If you run into an error about the inability to run the query because temporary sort space is exceeded, increase the temporary sort space as follows:

Ask the ESM administrator to edit the `/opt/arcsight/logger/data/mysql/my.cnf` file in the Manager installation. Edit the statement as follows:

```
sort_temp_limit = 1G
```

You may also want to further refine queries being used by reports and trends.

Using Queries in Query Viewers

You can use queries built for reports in *query viewers*, outside of the reporting paradigm. Query viewers provide a channel-style view of SQL query results but are not limited to events in terms of scope. They provide high-level summaries to monitor system health, reveal trends, and allow for drill-down and investigation of all types of resources across time. Query viewers are performance-tuned to work with trend tables rather than event tables, and so can return results much faster than active channels.

Query viewers include their own simple reporting option by which you can initiate a report on the query results from the query viewer.

For more about using query viewers, see ["Query Viewers" on page 340](#).

Building a Query



Caution: Do not exceed more than 10,000 resources in a group.

Where: Navigator > Resources > Reports > Queries tab

The high-level steps for creating a query are as follows:

1. Right-click a group (folder) and select **New Query**. This launches the Query Editor in the Inspect/Edit panel.



Note: As a general rule, create new content in the user's own folder.

2. Define [General Query Attributes](#). At a minimum, fill in the required values (red asterisks) on the General tab.
3. Define a schema for [Query Fields](#).
4. Create [Query Conditions](#).

5. Define Query Variables (optional).
6. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
7. Click **Apply** or **OK** to create the new query.



Note: Be sure to click **Apply** or **OK** frequently to save settings intermittently as you work through the above steps. Clicking **Apply** saves settings and leaves the Editor open. Clicking **OK** saves settings and closes the Editor for this query. If you do not apply or accept settings using these buttons, your settings are not saved.

Query Settings

Use the Query Editor to build a new query or edit an existing one. Query settings are defined on multiple sub-tabs.

Query settings include:

- ["General Query Attributes" below](#)
- ["Query Fields" on page 324](#)
- ["SELECT Query Fields" on page 325](#)
- ["GROUP BY Query Fields" on page 328](#)
- ["ORDER BY Query Fields" on page 331](#)

General Query Attributes

The following fields in the **Query** section are required attributes for creating queries.

General Query Attributes

Query Fields	Description
Name	Name for the query. Spaces and special characters are OK. This is an alias for the query that appears in pick lists in other editors.
Query on	From the drop-down menu, select one of the following data sources: • Event - Select Event if you want to create a report or view trends on event activity • Active List - Select Active List to query or view trends on list entries. Additionally select a Query Type. If you are creating a list query with asset-related conditions, see also "Example: Creating Asset-Related Conditions for Queries on Lists" on page 337. For more about active lists, see "List Authoring" on page 486. Cautions: ◦ For good query performance, query on case-sensitive lists only. ◦ If the active list contains list fields that reference resources (for example, a field that uses any Group variable function), and the query is used by a trend, that trend will not display the list of resource references. Rather, the trend will display only one list element. This limitation does not affect reports and query viewers. • Actor - Select Actor to query or view trends on actor information. (For more information on actors, see "Actors" on page 775.) • Asset - Select Asset if you want to report or view trends on statistics about the assets on your network, such as a list or count of assets categorized in a particular asset category, or the zone a particular asset is in at a particular time. (For more about assets, see "Modeling the Network" on page 119.) • Case - Select Case if you want to report or view trends on the status of cases, such as number of cases opened and resolved. (For more about cases, see "Case Management and Queries" on page 624.) • Notification - Select Notification if you want to report or view trends on the status of events sent out in the notification workflow, such as number of events in the Investigate stage. (For more about notifications, see "Managing Notifications" on page 216.) • Session List - Select Session List to query or view trends on session activity. If you are creating a list query with asset-related conditions, see also "Example: Creating Asset-Related Conditions for Queries on Lists" on page 337. For more about session lists, see "Managing Session Lists" on page 500. Caution: For good query performance, query on case-sensitive lists. only. • Trend - Select Trend if you want to report or maintain trend information on the data gathered in another trend. For instructions about how to build a trend, see "Building Trends" on page 467.
Query On Resource	Available for queries on active and session lists. Select a list from the drop-down panel.
Query Type	Available for queries on active lists. Select one: • Snapshot - Select Snapshot if you want the query to return values from the active list with no historical baseline. • Interval - Select Interval if you want to view values within a specified period.

General Query Attributes, continued


Query Fields	Description
Start Time	This field only appears if you are querying on an interval active list, event, or trend. Enter values depending on the data source you selected: • Active List, Interval type - Specify the starting point for the data gathering from the specified active list. • Event - Specify the starting point for the data gathering from the events database. Event data is generally kept unarchived for 30 days by default, so specify a start time within that time frame. • Trend - Specify the starting point for the data gathering from the trends database. Be sure to specify a period within the lifecycle of the trend; otherwise, the query returns an empty result set. Tip: If the query is used as a base query in a trend, the trend start time overwrites the start time set here. See "Trend Parameters" on page 476.
End Time	This field only appears if you are querying on an interval active list, event, or trend. Enter an end time depending on the type of source data you selected: • Active List, Interval type - Specify the ending point for the data gathering from the specified active list. • Event - Specify the ending point for the data gathering that is some time after the starting point. Keep in mind that large time spans can mean large amounts of data, which can affect system performance. • Trend - Specify the end point for the data gathering that is some time after the starting point. Tip: If the query is used as a base query in a trend, the trend end time overwrites the end time set here. See "Trend Parameters" on page 476.
Use as Timestamp	This field only appears if you are querying on an interval active list, event, or trend. This field indicates which value to use as the timestamp for the report itself. This value helps with sorting and scheduling. The following options are available for queries on events and trends: • End Time - Select End Time if you want to use the event or trend end-time you specified in the End Time field. The timestamp reflects the event end time. If you are querying on a trend, select this option. • Manager Receipt Time - Select Manager Receipt Time to use the time the event was received at the Manager. If you are querying on a trend, this is probably not an appropriate option to choose because in that case, Manager Receipt Time would indicate when the trend is run, rather than when events are received by the Manager. The following options are available for interval queries on active lists: • Date-based field on the active list - This is the default, if such field exists in the active list. • Creation Time - When the list was first populated (created) • Last Modified Time - When the list was last updated

General Query Attributes, continued

Query Fields	Description
Row Limit	Set the row limit for the data table. The default is 10000 rows. Tips: - The row limit you set here determines the row limit for reports using this query. Consider how row limit will affect report readability. For example, if you have a simple chart with just the X- and Y- axes, you might want a maximum of 20 rows for a single-page chart. For stacked charts, your data points still correspond to the row limits but two or more will be on the same column. See also "Selecting Data for the Z-Axis on a Chart (Optional)" on page 405 for additional information. - If the query is used as a base query in a trend, the trend row limit overwrites the row limit set here. See "Trend Parameters" on page 476.
Distinct Rows	This setting means only unique (distinct) rows appear in the results. For example, if you checked this box and there are duplicate returned rows, only one of them is shown.
Database Hint	This option does not apply to CORR-Engine.

Example:

The following example shows a query definition named *VPN Logins Outcome - Hourly*. Each time you run this query, it returns VPN login attempts over a one-day period each time based on Start Time (\$Now - 1d) and End Time (\$Now).



Tip: Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

Query Fields

The Query **Fields** tab contains the following main options with which to define query data and structure:

- [SELECT Query Fields](#)
- [ORDER BY Query Fields](#)
- [GROUP BY Query Fields](#)

Drag-and-drop on Query Structure panels:

You can drag-and-drop items between options (for example, to group by Category Outcome, drag it from SELECT to GROUP BY. It stays in SELECT but is also used to GROUP BY).

Search shortcuts:


- Enter part of the field name to find (for example, Name) in the Search box.
- Use the up/down arrow keys to jump to each instance of “Name” in the available fields.
- When you find the field name you want, press Enter to add it to the condition statement under the selected section (SELECT, GROUP BY, or ORDER BY)
- **Ctrl+F** brings back the Search box in focus if it is hidden

Common Conditions Editor (CCE):

The Query Editor, like other resource editors, uses the CCE for building conditional statements (query structure). For more tips on using the CCE, see "[Common Conditions Editor \(CCE\)](#)" on [page 887](#).

SELECT Query Fields

Click **Add SELECT columns** to select the data for the query. Data selected enters one big bucket, and any functions set for any of the data fields is performed on the entire bucket of data.



✓ **Tip:** Fields in shown in italics on the Data Options panel are derived, referenced, or side table fields (rather than "hard event data" in the main database tables). See also, ["Data Fields" on page 910](#) and ["Variables" on page 1064](#).


Query Structure (SELECT)

The Query Structure section at the top provides a summary of the fields selected in the SELECT section at the bottom. If you add GROUP BY or ORDER BY settings, these show up here also.

You can select from **Fields and Global Variables**, **Field Sets**, or **Local Variables** as data to build the query. Choosing a field set limits the fields shown to the selected field set.

- Click a field or variable to select it.
- Click again to deselect it (remove the checkmark).

- To edit a field or variable that you already have set as a query condition (showing under SELECT), simply double-click it or select it (click once) and click the Edit button () in the toolbar. (For example, you might want to edit the query by adding a function to it, as described in [SELECT Query Fields](#).)
- To duplicate a field or variable that you already have under SELECT, select it (click once) then click the Duplicate Column button () in the toolbar.
- To move column up or down, select it and click the up or down arrow in the toolbar.



You can also select a condition item and right-click to get the various Edit options (**Edit**, **Copy**, **Delete**, **Duplicate**, and so forth).

Applying Functions to SELECT Columns

Optionally, you can specify an aggregate function on a particular column of data, such as a line item count, or in the case of numeric data, a sum or average

If the query is not grouped by one or more columns, then aggregate functions added here are applied to the whole result set.

If the query is grouped by one or more columns, then the aggregate function is performed on each group individually.

Adding a function adds a data field to the query schema that provides the results of the function, which can later be displayed in a report.

To specify a function for column data, double-click a field or variable in the top pane under SELECT and select a Function from the drop-down menu to apply to the column data.

The available functions are:

- **COUNT** - Count the number of line items returned in this column.



Note: The Count function is a simple count of all events. It takes into consideration the aggregated event count and counts each event in an aggregated event individually. For example, if an event has an aggregated event count of 5, the Count function counts this event as equivalent to 5 events (with an aggregated event count of 1 each). Take this into account when comparing the number of rows in a report with the “grand total” count based on the Count function.

- **SUM** - Add all numerical data in a column, such as aggregated event count.

- **AVERAGE** - Calculate the average of all numerical data in a column, such as aggregated event count.
- **MAX** - Calculate the top values of the items returned in this column.
- **MIN** - Calculate the lowest values of the items returned in this column.
- Standard Deviation (**STDDEV**) - Calculate the variation from the “average” (mean) for this column. (Square root of the variance.)
- **VARIANCE** - Calculate the amount of variation within the values returned for this column.
- **GROUP_CONCAT** - Create a comma-separated list of the aggregated items.

Select **Unique** to apply the function only to unique values in the column. For example, the target address column may have 50 items in it, but only three are unique. To get a count of unique target addresses, check the Unique box.


Click the green checkmark button () to add the function.

To remove a function from a field, select the field, change the function selection to None, and click the green checkmark button again.

To cancel a modification to a function, click the () button or simply click elsewhere on the UI (off of the Function menu.)

GROUP BY Query Fields

Click **Add GROUP By** to divide query results into separate buckets. For example, you could do a “group by” if you are interested in sorting items by timestamp, such as logins between 3 and 5 p.m. Functions on **GROUP BY** data apply to timestamp-based fields only.



After fields are added to Group By, the Conditions tab's subtab, Group Conditions, becomes enabled.



Tip: Fields shown in italics on the Data Options panel are derived or referenced fields (rather than "hard event data" in the main database tables). See also, ["Data Fields" on page 910](#) and ["Variables" on page 1064](#).

Query Structure (GROUP BY)


The Query Structure section at the top provides a summary of the fields selected in the GROUP BY section at the bottom. SELECT and ORDER BY settings show up here also.

Adding and editing fields and variables to order by works similarly to adding them for SELECT. See ["Query Structure \(SELECT\)" on page 326](#).

Applying Time-Based Functions to GROUP BY Columns

You can specify a time-based function on the group by column of data. Time-based functions apply only to time-based fields, such as event end time.

To specify a function for GROUP BY column data, double-click a field or variable in the top pane under "GROUP BY" and select one of the available time-based functions (from the drop-down menu) to apply to the column data.



Functions on items under GROUP BY create a separate bucket of data for each time function specified.


To specify a function for column data, select a data field in the Query Columns section then select a Function (from the drop-down menu) to apply to the column data:

- **Second** - Creates a new bucket for all events that occur in the same second.
- **Minute** - Creates a new bucket for all events that occur in the same 60-second period.
- **Hour** - Creates a new bucket for all events that occur in the same 60-minute period.
- **Day** - Creates a new bucket for all events that occur in the same 24-hour period.

- **DayofWeek** - Creates a new bucket for all events that occur on the different days of the week, such as Monday, Tuesday, and Wednesday.
- **DayofMonth** - Creates a new bucket for all events that occur on various days of the month, such as the first, second, and third.
- **Week** - Creates a new bucket for all events that occur in a week.
- **Month** -Creates a new bucket for all events that occur in a month.
- **Year** - Creates a new bucket for all events that occur in a year.
- **Quarter** - Creates a new bucket for all events that occur in a quarter.

ORDER BY Query Fields

Click **Add ORDER BY** columns to specify the order in which you want the data in your buckets sorted. For example, you might use ORDER BY if you were interested in the numeric value of the items in your bucket such as the top 10 logins.



Tip: Fields in shown in italics on the Data Options panel are derived, referenced, or side table fields (rather than "hard event data" in the main database tables). See also, ["Data Fields" on page 910](#) and ["Variables" on page 1064](#).

Query Structure (ORDER BY)

The ORDER BY columns can be different than the ones you chose for the query data under SELECT. Also, you can apply functions to these columns.


Adding and editing fields and variables to order by works similarly to adding them for SELECT. See ["SELECT Query Fields" on page 325](#).

Applying a Column Function to Order By

Optionally, you can specify an aggregate function on a particular column of data to group by, such as a line item count, or in the case of numeric data, a sum or average.

You apply a function to ORDER BY columns the same as you do to a SELECT column, and the same functions are available depending on the fields or variables chosen. See ["SELECT Query Fields" on page 325](#).

To specify a function for column data, double-click a field or variable in the top pane under "ORDER BY" and select a Function (from the drop-down menu) to apply to the column data.




Sort Order

Under ORDER BY you can also set the sort order on the fields/columns. By default, the sort order is ascending (**ASC**). You can change it to descending (**DESC**).

Query Conditions

Optionally, you can create conditions on individual fields or on groups as part of the query. You can add filters, and conditions based on assets, vulnerabilities, and active lists. Query conditions represent the WHERE clause of the query.

Use the [Common Conditions Editor \(CCE\)](#) within the query editor to create query conditions as described in this section.



Tip: The Common Conditions Editor is used throughout the ArcSight Console for various resources. In addition to the topics that follow on defining conditions for a report query, see also ["Common Conditions Editor \(CCE\)" on page 887](#), ["Conditional Statements" on page 903](#), ["Conditions" on page 904](#), and ["Logical Operators" on page 1006](#).

Creating Conditions on a Field

See also ["Logical Operators" on page 1006](#), ["Condition Tree Command Buttons" on page 890](#), ["Condition Tree Context Menu Commands" on page 892](#), ["Common Conditions Editor \(CCE\)" on page 887](#), and ["Adding Conditions" on page 895](#).

1. Click the **Conditions** tab and select data fields from the data fields table (by default, fields from Common Condition Editor are shown). The fields you select are used to build condition statements in the display area at the top of the Edit sub-tab.

The data fields table displays a **Name**, **Operator**, and **Condition** column. These three columns are combined to create *<data field> <logic operator> <data field value>* condition statements. For example, if monitoring a Cisco Router, you could define a condition statement to specify `Device Product = Cisco Router: Device Product` as the data field, `equals (=)` as the logic operator, and `Cisco Router` as the data field value.

2. In the Op column, double-click the cell and select a logic operator from the drop-down menu.

3. In the Condition column, enter a data field value or double-click the cell and select a value from the drop-down menu. Press **Enter** to add the condition to the statement.
4. Repeat this process to add more statements to the condition.
5. Click **Apply** or **OK** to save your changes and create the condition.

Creating a Group Condition

Creating a group condition is similar to creating a normal condition, except you pick an aggregate function to perform on a field.

You would use it if, for example, to group by event name and when you want to get only the events with more than 100 occurrences in the query. In this case, you would add a `Count()` aggregate function to the `eventId` field, for example, `count(eventId) > 100` to eliminate the events that have occurred less than 100 times.

The Group Conditions subtab on a query's Conditions tab is enabled if the Group By section in your SELECT statement contains fields. On this subtab, the fields table at the bottom of the panel has the **Function** column in addition to the standard columns **Name**, **Op**, and **Condition**.

Tips on Creating Conditions

- Drop-down menus appear if the selected data field has a set of value options.
- For example, if the Category Behavior data field is selected, a drop-down menu appears with the value options of `/Access`, `/Access/Start`, `Access/Stop/` and so on. One of the choices in this menu is `/Authentication/Verify`, which is the condition we selected for Category Behavior in our example condition.
- For date and time data fields, such as Detect Time, you can type an actual date value, such as `10/12/2017 8:54:00 AM`, or you can use special Time variables.
- The condition statement appears as a branch under the logical operator.
- To add a condition to an event field, click in its condition box and click the ellipses icon.
- To activate all operands on the top, select an item in the editor view, as shown above.

Query Variables

Variables are run-time information derived from the source data (event, asset, case, notification, or trend, depending on the schema) that can be used in the query wherever normal fields can be used. These variables are local only to the resource where they are created (in this case, a query).



Tip: You can create global variables that can be shared across resources. For information, see [Global Variables](#).

To set a local variable:

1. Click the query's **Variables** tab.
2. Click **Add** to open the Variables dialog.
3. The Variables dialog displays different values depending on the function you choose. In the Variables dialog, enter the following values and click **OK**.

Options	Description
Name	Enter a name for the variable. This is the alias that appears in the Conditions editor when you can use the variable. Spaces and special characters are OK. Note: Ensure that the name is unique across resources. Local variables cannot share names.
Function	From the drop-down menu, select a function. For a description of each function, click Help in the lower right corner.
Arguments	The arguments section contains a series of fields where you set the parameters for the variable. The available fields vary with the function you select.
Preview	The preview area provides an interface where you can enter values for the key variable fields so you can verify that the parameters you specified return the expected results. Enter test values and click Calculate .

Editing a Query

1. Navigate to **Reports** in the Navigator panel, select the **Queries** tab, and select the query you want to modify.
2. Double-click the query, or right-click and select **Edit Query** from the context menu. This launches the Query Editor in the Inspect/Edit panel, and shows the definition for the selected query.
3. Edit the query definition as needed. See "[Query Settings](#)" on [page 321](#) for details on query attributes.



Caution: If you are editing a trend query

If the query is used in a trend, the query and associated schema referenced in the trend are set at the time the trend was created. After the trend is created, you can add columns to the base query, but columns added to the query after the trend is created are not used by the trend. You can remove columns from the base query that are not used by the trend.

However, if you want a trend to have the added or removed columns in the trend query, create another trend and select the modified query.

To view the full conditions for the **MatchesFilter** operator, click the **Summary** tab and then click the **Expand Filter** button to display the filter conditions for debugging.

Note that in this case, the display of the **MatchesFilter** full logic does not display the sub-filter of the matched filter. Full logic is displayed only for the first level of matched filter conditions.

4. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
5. Click **Apply** or **OK** to save your changes. (Click **Cancel** to exit the Query editor without saving changes.)

Example: Creating Asset-Related Conditions for Queries on Lists

This example applies to queries on active and session lists. A list is typically populated automatically by one or more rule actions (see ["Rule Actions Reference" on page 539](#)). You can then query those lists to create reports.

For queries on lists, the Assets button on the query's Condition tab is disabled. This therefore limits your ability to add an asset-related condition by means of a button. You can, however, use the Common Conditions Editor to create your query condition statements in queries on lists containing assets.

This example uses a query on a list that contains assets. The goal of this example is to use a list to collect asset information, and you would regularly query that list for this information. Specifically, you are looking for assets that belong to a specific category.

Model your network:

Assets are part of your network model. For complete information, see ["Modeling the Network" on page 119](#).

Create your list:

Create a field-based active or session list containing assets fields. Details are in ["List Authoring" on page 486](#).

Create your query:

This query is interested in looking at the list and getting assets that belong to a particular asset category. Details for setting general query attributes are in ["General Query Attributes" on page 321](#).

Add your query conditions:

For our purposes, we want to look for assets that belong to a category.

1. On the Conditions tab's Common Conditions Editor, locate the asset field of interest, for example, Asset ID. Use the **InGroup** operator for the field.
2. Click the ellipsis in the Condition column to display the Advanced Editor for the asset field.

The Advanced Editor provides the following options for the asset field:


Asset Options	Action
Asset	Expand the asset group nodes and select an asset group.
Zone	Expand zone nodes and select a zone.
AssetCategory	Expand the asset category nodes and select an asset category.




Note: The Asset options listed above are not available to all fields.

3. Specify your asset attribute of interest and browse to the desired resource URI to complete your condition statement.

For example:



In our example, we choose **AssetCategory**. We are interested in the Criticality category that is equal to High. A list of asset category groups is displayed, from which you can further expand to select the category you want. The query's condition statement is then updated according to your selections, for example:



The query returns Asset IDs belonging to the asset category specified in the condition statement.

Chapter 12: Query Viewers

Query viewers are used to define and run SQL queries on other resources, including trends, assets, cases, connectors, events, and so forth. Each query viewer contains a SQL query along with other logic for establishing and comparing baseline results, analyzing historical data to find patterns in network activity, and performing drilldown investigation on a particular aspect of the results. The query viewer you create displays all the fields specified in the query you select (or create) for the query viewer.

Pre-Built and Custom Query Viewers

The Manager to which your Console is connected has pre-built query viewers available for use. At a minimum, you have access to standard content query viewers that ship with ArcSight. You might also have access to custom query viewers provided by content developers for your organization.

Standard Content

ArcSight comes with a set of pre-built query viewers that address common network monitoring and trend analysis scenarios.

To access the standard content query viewers:

1. In the Navigator panel select **Query Viewers**.
2. Click to expand the list.
3. Locate and select **Query Viewers/Shared/All Query Viewers**.

Refer to the [ArcSight Administration and ArcSight System Standard Content Guide](#) for the descriptions of the pre-built queries in this query viewer group.

For information on how to run and use any pre-built query viewer, see "[Viewing Query Viewer Results](#)" on page 362, "[Generating Reports from Query Viewers](#)" on page 374, and "[Defining and Using Baselines](#)" on page 349.

Custom Query Viewers

When administrators or content developers at your organization create custom query viewers, they have the option of sharing these with other administrators and users. So, depending on your role and user permissions, you might have access to:

- Query viewers that ship with ArcSight
- Custom-built query viewers that other administrators have shared with other ESM users
- Your own custom-built query viewers

For information on how to create your own custom query viewers, see ["Managing Query Viewers" below](#).

Customizing Query Viewers as Needed

You can modify the provided query viewers as needed to get the data you want. Customizing an existing query viewer can range from hiding or showing data fields, changing the sort order inherited from the base query, to adding variables and modifying key fields. These kinds of changes do not affect the base query, only the query viewer.

Once a query viewer is defined to reference a particular base query, that query viewer cannot be changed. If you want to reference a different base query, create a query viewer. For additional information, see ["Base Queries" on page 876](#).

inActiveList Conditions for Queries

In a query, you can define an `inActiveList` condition and map multi-valued attributes to single-valued active list fields. For example, you have an active list that keeps track of actor roles, where one of the role values can be Normal, Restricted, and Privileged. You can test if an actor has one of these roles through the `inActiveList` condition. In this scenario, your list has a field called `RoleName`. You map the actor's role name attribute to this field. Keep in mind that an actor's `RoleName` attribute is multi-valued because an actor can have multiple roles. Through the `inActiveList` condition, your query will check if one of the actor's roles is, for example, Privileged.

Querying an active list of assets require extra steps, as explained in ["Example: Creating Asset-Related Conditions for Queries on Lists" on page 337](#).

Managing Query Viewers

Query viewers provide a shortcut alternative to running SQL queries as a part of reporting. Keep in mind that query viewers use base queries, so a first step in creating a query viewer is deciding what SQL query you want to use. If you can't find one that does what you want, you'll need to create one first, then use it as the base query for a new query viewer.

 **Caution:** Do not exceed more than 10,000 resources in a group.

The high-level steps for creating or editing a query viewer are as follows:

1. Identify your questions and what information you are looking for. (For example, "What types of actions represent the highest volume of events on my network during various times of day?")
2. Based on the question you want answered, decide what kind of query you need and determine whether it is available or you have to create it.

If you have not created a query for this query viewer, see ["Building a Query" on page 320](#).

3. Select **Query Viewers** in the Navigator.
4. Right-click a group (folder) and select **New Query Viewer**.



Note: As a general rule, create new content in your own user's own folder.

To edit a query viewer, right-click and choose **Edit Query Viewer**.

5. In the edit panel, define general attributes for the query viewer as described in ["Query Viewer Attributes" on the next page](#). At a minimum, fill in the required values (red asterisks) on the **Attributes** tab (query viewer name and base query to use).



Note: You cannot change the base query after the query viewer is saved. If you want to use a different base query, create another query viewer for it.

6. Select the **Fields** to display for the query viewer as described in ["Query Viewer Fields" on page 345](#). (Fields are inherited from those available in the base query.)
7. Define **Variables** for use in the query viewer as described in ["Query Viewer Variables" on page 348](#) (optional).
8. Specify any **Drilldowns** you want to include with the query viewer as described in ["Managing Drilldowns from Query Viewers" on page 356](#) (optional).
9. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
10. Click **Apply** or **OK** to create the new query.



Note: Click **Apply** or **OK** frequently to save settings periodically as you work through the above steps. Clicking **Apply** saves settings and leaves the Editor open. Clicking **OK** saves settings and closes the Editor for this query. If you do not apply or accept settings using one of these buttons, your settings are not saved.



Tip: To edit a query viewer for which results are currently displayed in the Viewer, click the Edit Query Viewer button that looks like a pencil on the lower right of the Viewer.

The results display for the query viewer you want to edit must have focus (that is, be on top) in the Viewer.

Query Viewer Settings

Use the Query Viewer Editor to build a new query viewer or edit an existing one. Query viewer settings are defined on multiple sub-tabs.



Tip: You can access the editors for multiple query viewers simultaneously.

- To access the editor for a query viewer, follow the first steps in ["Managing Query Viewers" on page 341.](#))
- If you want to edit more than one query viewer at a time, choose **Edit > Preferences** from the Console menu, then click **Global Options**. On the Global Options panel, check **Allow multiple editors of the same type**, then click **OK** to save the change and close the Preferences dialog. For more on setting Console preferences, see ["Changing Global Options" on page 49.](#)

Settings include:

- ["Query Viewer Attributes" below](#)
- ["Query Viewer Fields" on page 345](#)
- ["Query Viewer Variables" on page 348](#)

Query Viewer Attributes

The following fields in the **Query Viewer** section are attributes to specify when creating a new query viewer.

Query Viewer Attributes

Query Fields	Description
Name	Required: Enter a name for the query viewer. Spaces and special characters are allowed.
Query	Required: For first-time query viewer configuration, specify the base query used in this query viewer. • Select the query using the "Select a Query" drop-down menu. The arrow displays the Queries resource tree. • Alternatively, if the resource tree has too many subgroups to traverse and you know the query by name, see "Using the Advanced Selector While Editing Resources" on page 88. Note: If you are editing an existing query viewer, the Query field is not editable. If you want to use a different base query, create another query viewer.

Query Viewer Attributes, continued

Query Fields	Description
Refresh Data After	Set an amount of time (in minutes or hours) after which the query viewer automatically runs again and shows new data based on that most recent run. The query viewer is regularly refreshed based on the specified refresh time period. The default for this setting is after every 15 minutes. To change this default: 1. Click the field to activate the settings. 2. In the left-hand field, enter a numeral, and in the right-hand drop-down menu, select minutes or hours.
Query Time Out	Define a time out limit in which the query must return results. If the query does not complete and sends no results within the specified time out period, the Manager stops running the query. By default, the time out is 300 seconds (5 minutes). If you do not specify a Query Time Out in the Attributes tab, this time out of 5 minutes applies, even if the Query Time Out field displays None. Setting a time out limit is good practice especially if the event rate (events per second or <i>EPS</i>) is unusually high, start/end time range is large, or the query is complex. Time outs can help guard against infinite or long running queries that impact system performance. Although this is less of an issue with query viewers since they are designed to minimize impact on system performance, this can still be an issue in some scenarios. Setting time outs can be a useful troubleshooting technique for new queries, or existing queries in new scenarios, for example where event counts spike higher.
Default View	The Default View attribute determines how the result data are displayed when you double-click the query viewer to open the results in the Viewer panel. Define the default (double-click) view format for this query viewer. The choices are to show data as: • Table (this is the default) • Pie chart • Bar chart Double-clicking a query viewer in the Navigator displays result data in the format set here. If you choose Pie Chart or Bar Chart as the default view format, choose fields to use for the Values Column (to plot the y axis points on a bar chart or slice sizes on a pie chart) and Points Labels column (to plot the x axis labels on a bar chart or slice labels on a pie chart). The Values Column and Points Labels are also described in "Viewing Query Viewer Results" on page 362.
Values Column	The Values field applies to bar charts and pie charts. This setting provides fields in the query result that contain data types. The value chosen is used as the numbers by which to plot the vertical y axis points on a bar chart or the slice sizes on a pie chart. Values typically represent an unknown set of values, like a count. A common example of numeric data appropriate for values is a time like HourOfDay or a count like Count (Event ID).

Query Viewer Attributes, continued

Query Fields	Description
Point Labels Column	The Point Labels field applies to bar charts and pie charts. This setting provides fields in the query result that contain non-numeric data types. The point labels are used to plot the horizontal x axis labels on a bar chart or the slice labels on a pie chart. Examples of non-numeric data types appropriate for point labels are timestamps, strings such as are used for event names, and different types of addresses such as IP or MAC addresses. Point labels are typically a known set of limited values (like hours in a day denoted by timestamps).
Setting the following attributes (Start Time, End Time, or Row Limit) in the Query Viewer overrides these settings in the base query. (See Query about defining the base query in the <i>Query</i> attribute.)	
Start Time	Specifies the starting point for the data gathering. A drop-down menu provides values to select based on Velocity Templates (such as \$Now, \$Now - 1d, and so on). You can also provide a timestamp such as: 27 Jul 2017 16:00:00 PDT. For more on timestamps and timestamp variables, see "Timestamps" on page 1059, "Timestamp Variables" on page 1060, and "Variables" on page 1064.
End Time	Specifies an end point for the data gathering. A drop-down menu provides values to select based on Velocity Templates (such as \$Now, \$Now - 1d, and so on). You can also provide a timestamp such as: 28 Jul 2017 16:00:00 PDT. For more on timestamps and velocity references, see "Timestamps" on page 1059, "Timestamp Variables" on page 1060, and "Variables" on page 1064.
Row Limit	Set the row limit for the data table. The default for all new base queries is the maximum allowable, which is 10,000 rows. If the default is not changed in the base query, and no limit is specified here in the query viewer, the result shows up 10,000 rows of data.

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

Query Viewer Fields

To define the data display, click the query viewer **Fields** tab.

Inspect/Edit

Event Inspector Query Viewer: Event Details

Attributes Fields Local Variables Drilldowns Notes

Data Fields

☒	Name	Display Name	Key
☒	End Time	End Time	☐
☒	Name	Name	☐
☒	Priority	Priority	☐
☒	Attacker Address	Attacker Address	☐
☒	Target Address	Target Address	☐
☒	Device Address	Device Address	☐
☒	Device Product	Device Product	☐
☒	Device Zone Name	Device Zone Name	☐

Sort Options

+ Add... Remove

Column	Sort Order

Baselines

Remove

Description

OK Cancel Apply Help

The data fields shown on this tab are inherited from the base query. When a query viewer is first created, the data fields are shown here with the same settings they inherited from the base query for Use and Key fields. So, initially all fields are enabled for Use and fields that are grouped by columns in the base query show as Key fields here.

You have the option of overriding the base query settings for Use and Key settings on inherited data fields in the query viewer. (Settings here do not affect the base query.) You can override these settings when you first create the query viewer, or when you edit it later.

Select (check) **Use** for fields to display in the query viewer results. Fields not selected to Use do not show up in the query results.

Optionally, you can select one or more fields to use as **Key** fields. Key fields are columns that can be used to uniquely identify a role in the query. Only the fields selected as keys are used when doing baseline comparisons.

Name	Alias	Use	Key
TimeStamp	TimeStamp	☒	☐
Name	Name	☒	☒
Count(Event ID)	Count(Event ID)	☒	☐

The query viewer displays results from these columns, showing them from left to right in the order specified. The above settings would result in a query viewer that shows Timestamp as the left-most column, followed by Name, and so forth. You can re-order the columns by selecting a row and clicking the up or down arrow to move it.

Sort Options

The query viewer inherits the sort options from the base query, but you can override those sort options here, without affecting the base query.

You can add data fields from the base query to sort the query results in the query viewer display.

Click **Add** () to get the list of available fields and select those you want to sort on.

Column	Sort Order
TimeStamp	Z-A 
Count(Event ID)	A-Z 

In the example above, the Timestamp is sorted from newest to oldest. Data with the newest Timestamp is at the top of the list. Data with the oldest Timestamp is at the bottom of the list. (This is indicated by the Z-A sort order and up arrow.) In a case where multiple rows have the same Timestamp, these are sorted by the Count(Event ID) from smallest to largest (as indicated by the A-Z sort order and down arrow).

You can change the priority of a column by selecting a column and clicking the up or down arrow to move it.

Note: It is possible to sort on fields that you choose not to display in the query result.

Suppose you decide to hide the timestamp and count (event ID) columns. In the query viewer Sort Options, you can still sort by Count (Event ID) and Timestamp.

The list of event names and results for this query viewer display in this multi-column sort order by timestamp and count (event ID), but those *columns* do not show up in the display.

Baselines

If any baselines have been set on results returned on this query viewer, those are listed in the Baselines area of the Fields tab.

Baselines are created on query results tables using the right-click popup option **Analyze in Channel >Add as baseline...** after a query runs. (See ["Defining and Using Baselines" on page 349.](#))

When a query has one or more baselines available, you can compare the current results of a table view with the baseline.

To remove baselines from the query viewer, click the **Fields** tab, select the baseline name, and click **Remove** ( Remove). Be sure to click **OK** or **Apply** on the Query Viewer Editor to save your changes.

If you remove the baselines from the query viewer definition, they are not available on the next run of the query viewer.

Query Viewer Variables

To add a local variable:

1. On the **Local Variables** tab, click **Add**.
2. Provide a name for the local variable.



Note: Ensure that the name is unique across resources. Local variables cannot share names.

3. Choose a function from the **Function** drop-down list.
For a description of available functions, see [Variables](#).
4. Provide other details as needed and click **OK**.

The local variable is available in the following views:

- As a field on the **Fields** tab in the query viewer editor definition (including the options to **Use** and use as a **Key** field).
- As a column in the query viewer result. If the query viewer result is displayed in the viewer when you add the variable, the variable shows up immediately as a column in the result.

For example, you can add a local variable using the Timestamp Function (such as `GetHour`, `GetDayOfWeek`, `GetDayOfMonth`, and so forth).



Note: You cannot promote a query viewer local variable to a global variable. Query viewers operate on queries, which have their own distinct schema for each instance. A local variable defined for a query viewer is only applicable to that query viewer.

Local variables defined for data from events, actors, cases, and assets can be promoted to a global variable.

For more information about global variables (which can be used in queries), see [Global Variables](#).

Deleting a Query Viewer

1. Navigate to **Query Viewers** in the Navigator panel, right-click the query viewer you want to delete, and select **Delete Query Viewer**.
A confirmation dialog is displayed.
2. Click **Delete** to confirm your choice and delete the query viewer.

Defining and Using Baselines

You can establish a particular set of query results as a **baseline** snapshot against which to compare the results of other runs of the same query. Comparing the results of the same query run at different times and in different contexts highlights the *deltas* (differences) and helps identify areas that vary significantly from normal. If spikes, dips, or other anomalies appear, you can compare them against the baseline.

You can define baselines and run comparisons with any query viewer that:

- Lends itself well to a table format display
- Includes one or more key fields by which to locate matching entries between the baseline and currently displayed information.

For example, suppose you have a query that returns the top 10 event counts by name and you want to compare it against some baseline. A reasonable comparison would be between similarly named events in both sets of data. In this case, the event name would be used as the key field.



Note: Following are important considerations on baselines.

- **Baselines are applicable only to table views of result data.** Baselines do not apply to graphical views such as pie charts, bar charts, and so on. You always have the option to view query data from any query viewer as a graphical chart or a table, but the baseline data is only accessible from the table view of the data.
- **Baselines require one or more key fields** by which to locate matches between the baseline and the displayed data. The key fields must be built into the query viewer to which you want to add a baseline.
- **Values for Key fields must be unique.** When adding baselines, make sure key fields in the query viewer have unique values. (See the **Fields** tab in the query viewer editor.) Also, check the query viewer start and end times (on the **Attributes** tab in the query viewer editor) to make sure the time frame over which the query runs makes sense.

You can add one or more baselines to a single query viewer, and delete them as needed.

Why Baselines are Useful

In addition to providing a way to compare result data from different query runs, baselines provide an efficient way to save, annotate, and retrieve data that might otherwise be too difficult to access in any meaningful way.

A baseline is preserved as a File resource that is associated with the query viewer. In the Navigator, choose **Files** and expand the **Attachments** folder to view the new baseline files you created.



Tip: The query viewer baseline files in the Files\Attachments folder appear along with other files in alphabetical order. For your convenience, in the Files resource, you may add a folder (add a group) to contain only your baselines. You can name it <yourname> Baselines and drag your own baseline files from Attachments to your own folder. The baseline files always remain in Attachments but a link is established from files in Attachments to the files in your own folder.

In **Query Viewers**, you can create, save, and use baselines to compare result data from the same query viewer run at different times and dates.

Baselines created in **Query Viewers** also show up under the **Files** resource in the Navigator. For convenience, create a folder to contain your baselines.

Baselines are automatically saved in the Attachments folder. For convenience, create another folder and drag your baselines there.

With Query Viewer baselines, you can:

- Retrieve the snapshot baseline data by running comparisons against it.
- Compare current result data against one or multiple baselines.
- Get meta-information about the baseline (such as when it was saved, by whom, and comments).
- Sort, show, or hide the baseline comparison columns.

- Maintain the baseline data as a Files resource baseline even if the original data is lost or is too performance-intensive to re-generate (for example, an aggregation query). (All baselines are automatically added as Files resources when they are created.)
- Add and remove baselines as needed, and edit some meta-information on baselines (for example, description comments).
- Use filters on the baseline (*delta*) columns. For example, you could filter on a baseline column to find where the current results differ from the baseline by more than some specified value.

Planning for Baseline Comparisons

Query viewer baselines might prove most useful if you take a little time to identify some goals for their use or questions you want answered, and then plan how to implement the baselines for those purposes. Here are some suggestions to start off with.

1. Establish questions or goals for baseline comparison monitoring and identify the type of data you want to evaluate.

For example, you might want to determine what type of event traffic is at its highest at different times of day or when network attacks tend to increase. Or, you might notice a spike in certain query viewer results (such as more logins from a particular user) and decide to compare the behavior against a sampling of results from subsequent or previous query runs.

2. Identify the query viewer (and associated query) appropriate to use. If the query viewer you need is not provided, you can develop it. See ["Managing Query Viewers" on page 341](#) for more information on this.

For example, if you want to monitor what type of event traffic is at its highest, you could establish a baseline for a query viewer that returns "Top Event Counts by Hour of Day." You could also use a query viewer *baseline* to take snapshots of event counts throughout the day, either for record-keeping or to explore and compare later.

3. Monitor results for your chosen query (by running the query viewer) to identify a "typical" or "normal" result set to use as a baseline.
4. Add (capture) the baseline from the typical/normal result set.
5. Monitor subsequent results for variation (spikes, dips) or time periods against which you want to compare with the baseline, and run baseline comparisons on these.

Adding a Baseline

You can only define baselines on numeric data because they are designed to show *deltas*, the difference or change between two values.

To add a baseline to a query viewer:

1. In the Navigator panel, choose the **Query Viewers** resource.
2. Select and run the query viewer (containing the query) for which you want to define a baseline.

To do this, right-click the query viewer and choose **View Data As > Table**.



Note: Baselines are applicable only to table views of result data.

The query viewer result is displayed in the Viewer.

3. Right-click anywhere in the results table in the Viewer, and select **Analyze in Channel > Add as baseline** to get the Add a baseline dialog.
4. Enter a name for the baseline, optional description, and click **OK** to add it.

This saves the current query result data as a named baseline for the selected query viewer, and makes it available for use (through **Analyze in Channel > Compare with...** against results from other runs of the same query viewer).

The baseline is shown on the query viewer's Fields tab.



Tip: If the query viewer editor is not currently displayed, double-click the same query viewer in the Navigator panel to open it in the editor. Click the query viewer editor **Fields** tab.

Comparing Displayed Results to a Baseline

After you establish a baseline for a query, you can compare subsequent results for the same query against the baseline.

Requirements:

- Baseline comparisons, like baselines, can only be derived from *table* views of the query viewer results. See ["Results in Table Format" on page 368](#).
- The query viewer you select for baseline comparison must have at least one baseline already added to it. Baselines are shown on the **Fields** tab of the Query Viewer editor.

To run a comparison:

1. If you do not already have a table view of the data you want to compare, right-click the query viewer you want to evaluate against a baseline, and choose **View Data as > Table** from the Navigator menu.

2. In the Viewer, right-click anywhere on the table view results and select **Analyze in Channel > Compare with: <SomeBaseline>**.

The comparison data is collected and added as a new column. You have the option of hiding or showing it in the table as needed.

3. Make your selections on the Select Columns table and click **OK**.



Tip: By the time the Select Columns dialog is displayed, the Baseline comparison is already available.

If you selected the comparison column, it is displayed on the table next to the original results for that column.

Events: Table		
Query: Events		
Last Update: 1 Aug 2016 16:40:14 PDT		
Filter: No Filter		
10 shown / 10 matches		
Name	COUNT(Events)	COUNT(Events) - (Typical Weekday Baseline)
Monitor Event	14071	-929
Top value count data monitor val...	10209	1209
ICMP PING	976	
ICMP PING *NIX	852	-57
ICMP PING BSDtype	852	-56
NETBIOS SMB-DS DCERPC NTLM...	678	74
NETBIOS SMB-DS Session Setup ...	317	99
AddToList: Success	238	0
ActiveList entry updated	233	0
Starting Trend Query	192	0

Note that differences between the current values and the baseline can be positive or negative, as shown in the example comparison above. A positive value in the baseline comparison indicates more events in your current sample, compared to the baseline. A negative value in the baseline comparison indicates fewer events in your current sample, compared to the baseline. If the baseline field for a row is null, this indicates that no baseline value was available for that key.



Tip: Here are more options available after comparing baselines.

- After running a baseline comparison, the right-click over Table **Analyze in Channel > Compare with <Baseline>** option for the baseline you just ran is grayed out (even if you chose not to immediately select any columns or clicked Cancel on the Select Columns dialog). This is because the baseline is already added.
- To show or hide more columns (including baseline columns), right-click the column header, choose **Show Column**, and check (enable) or un-check (disable) columns. See also ["Show or Hide Baseline Columns" below](#)

Show or Hide Baseline Columns

You can always show or hide columns, including baseline columns. To do this right-click anywhere in the table header (on any of the column titles), choose **Show Column >**

<SomeField>.

Name	Count(Event ID)	Count(Event ID) - (Typical Weekday Baseline) ▲
Sort Column	6990	6905
Remove Sort	5896	5811
Show Column	☒ TimeStamp	5734
Size Column to Fit	☒ Name	5726
Help	☒ Count(Event ID)	5619
Monitor Event	☐ Count(Event ID) - (Reference Baseline)	5545
Monitor Event	☒ Count(Event ID) - (Typical Weekday Baseline)	5375
Monitor Event	5460	5375

See also ["Column Sort, Display, and Edit Options" on page 370.](#)

Sort Baseline Data

You can perform an after-query sort on baseline comparison data by clicking the column headers. A pre-query sort for baseline data is not available. That is, there is no option to add a sort as a part of the baseline in the query viewer definition.


Viewer		
Events: Table		
Query: Events 10 shown / 10 matches		
Last Update: 1 Aug 2016 16:40:14 PDT		
Filter: No Filter		
Name	COUNT(Events)	COUNT(Events) - (Typical Weekday Baseline)
Monitor Event	14071	29
Top value count data monitor valu...	10209	19
ICMP PING	976	
ICMP PING BSDtype	852	66
ICMP PING *NIX	852	67
NETBIOS SMB-DS DCERPC NTLMs...	678	74
NETBIOS SMB-DS Session Setup A...	317	99
AddToList: Success	238	0
ActiveList entry updated	233	0
Starting Trend Query	192	0

See also ["Column Sort, Display, and Edit Options" on page 370.](#)

Filter Baseline Data

You can filter on the baseline comparison column the same way you would filter on any other column. Click the **Filter** in the query viewer header to bring up the Query Data Filter dialog. Enter your filter conditions and click **OK**. After the filter is applied, the query viewer automatically updates.

The Query Data Filter is based on the Common Conditions Editor (CCE). For information about using the CCE to define filters, see ["Common Conditions Editor \(CCE\)" on page 887.](#)




Removing a Baseline

Baselines, like the queries themselves, are associated with and contained in query viewers. To remove a baseline, you remove it from the list of baselines in the query viewer editor.

✓ **Tip:** Removing a baseline from a query viewer is different from hiding or showing a baseline column in a query result. If all you want to do is temporarily hide a baseline column in a results table, use the right-click "Show Column" option in the Viewer on the results table as described in ["Column Sort, Display, and Edit Options"](#) on page 370 in ["Results in Table Format"](#) on page 368.

To remove a baseline from a query viewer:

1. In the Navigator panel, right-click the query viewer containing the baseline you want to remove and select **Edit Query Viewer**.
This opens the editor for the query viewer in the Inspect/Edit panel.
2. In the editor, click the **Fields** tab.
3. Under **Baselines**, select the baseline you want to remove and click **Remove** ().
4. Click **Apply** to save your changes to the query viewer, or click **OK** to save changes and close the editor.

Note that there is no confirm dialog for this Remove baseline action, but if you do not want to save your changes, click **Cancel** and the baseline is not removed.

Managing Drilldowns from Query Viewers

Drilldowns provide the ability to investigate details about resources related to what is displayed by query viewers or data monitors. You can get more focused views on particular aspects of a single item, such as an asset, case, event, and so on, in the query result.

You can configure query viewers and data monitors to drill down to one or a combination of the following resources:

- Active channels
- Dashboards
- Query viewers
- Reports

Each drilldown type has its own options. After you have added one or more drilldowns, Console users can select one by right-clicking on the result and selecting **Drilldown > [drilldown name]** from the context menu.



Note: In a Custom View Dashboard and on the ArcSight Command Center, only drilldowns to dashboards are supported.

Adding a Drilldown

The drilldowns are initially displayed in the order they were created. The first drilldown is automatically the default.

To add a drilldown from a query viewer:

1. Access the **Drilldowns** tab in one of two ways:
 - Right-click on the query viewer or data monitor results in a dashboard and select **Drilldowns/Edit Drilldowns** to open the editor to the **Drilldowns** tab.

Or

 - Right-click on a query viewer or data monitor in the Navigator panel and select the **Edit** option, then select the **Drilldowns** tab.
2. Click **Add** () to open the Add Drilldown panel.
3. In the Destination field, select a resource type, for example, Dashboards.

The screenshot shows the 'Add Drilldown' dialog box. The 'Destination' dropdown menu is open, displaying a list of options: 'Dashboards', 'Active Channels', 'Dashboards', 'Query Viewers', and 'Reports'. The first 'Dashboards' option is highlighted. The dialog also contains a 'Menu Label' field with the instruction 'Give your drilldown a descriptive name.', a 'Description' field with the instruction 'Add more descriptive information about this drilldown.', and four buttons at the bottom: 'Next', 'Finish', 'Cancel', and 'Help'.


Then choose the corresponding specific resource, for example, My_Dashboard.

4. Enter a menu label (defaults to the specific resource's name). This label will represent this drilldown when the user right-clicks and selects Drilldowns on the Viewer panel.
5. Enter an optional description containing useful information about the drilldown.
6. Set the remaining options based on your destination resource:


Options for the Drilldown's Resource Destinations

If resource type is ...	Follow these steps ...
Active Channels	For an active channel destination, the settings in the Channel Display Options tab are not required; you may click Finish . If you want to set display options: - Select a field set from the drop-down list and click OK. - Change the Sort By field from the drop-down list and the sort order. - Click Finish.
Dashboards	Click Finish . You are done.

Options for the Drilldown's Resource Destinations, continued

If resource type is ...	Follow these steps ...
Query Viewers	For a query viewer destination, field mapping is required: - On the Field Mapping tab, click Add to display a dropdown list of source fields. You must define at least one field map. The source fields are from the source query viewer (the one you are drilling down <i>from</i>). The mapping condition is always set to =. - Under the Destination Field column, select a field from the destination query viewer (the one you are drilling down <i>to</i>).  The Drilldown definition shown in the example maps the source query viewer/data monitor “Name” column to the target query viewer/data monitor “Name” column. This constructs the following drilldown filter: <pre><target>.Name = <source>.Name</pre> where <source>.Name is replaced by the actual value from the source query viewer/data monitor row. If there are no eligible field mappings, you cannot complete the drilldown definition; the Finish button is disabled. You can add or remove field mappings, but your choices are limited to the columns already provided in the query viewer. - On the Display tab, you can choose to show (check) or hide (uncheck) the data fields in the drilldown result. - On the Sort tab, you can click Add to select the columns to specify the sort order of the resulting values. For each added column, change the sort order to ascending (the default) or descending. - Click Finish.

Options for the Drilldown's Resource Destinations, continued

If resource type is ...	Follow these steps ...
Reports	For a report destination, the settings in the Report Display Options tab are not required. To use the parameters set for the report, click Finish. If you want to change the drilldown's display options: - Click Add to display a list of the destination report's custom parameters, then select a parameter. - Under the Value column, select the field whose value will be used for the parameter.  - Click Finish.

7. Repeat the process to add multiple drilldowns as required.

The drilldowns you added will be available for selection when you view the data monitor or query viewer results. From those resources, the drilldowns are displayed for selection in the order they were created. The first drilldown is automatically the default drilldown of choice.

Tips on drilldown definitions:

- If there is only one drilldown, this is the default drilldown for that resource. If there are multiple drilldowns, the first drilldown is the default. You can change the order on the Drilldowns tab.
- When you run the query viewer results or view a data monitor, right-click, and select **Drilldown**, the selection list displays the list of drilldowns defined for that resource. The default drilldown is at the top of the list, and the remaining drilldowns are displayed in the sequence as they appear on the data monitor or query viewer's Drilldowns tab.
- You can define drilldowns for multiple fields of different data types. For example, you can define a drilldown to return a combination of event name and IP address. The first step would be to define a base query viewer to return these fields in a result, and then, as a next step, add a drilldown and select that query viewer to use as the “Drill down to” query viewer.
- You cannot define drilldowns to go to fields that are SQL functions.

Editing a Drilldown

To edit a drilldown:

1. Open the editor for the query viewer or data monitor you want to edit.
2. Click the **Drilldowns** tab.
3. Select the drilldown you want to edit and click **Edit** .

The drilldown dialog for this drilldown is displayed. Change the fields and options as described in ["Adding a Drilldown" on page 356](#).



Note: You can also edit the drill down from the query viewer or data monitor results. Right-click and select **Drilldown > Edit Drilldowns**. Selecting this command opens the editor for the query viewer or data monitor at the Drilldowns tab.

Changing the Default Drilldown

When you run the query viewer results or view a data monitor, right-click, and select **Drilldown**, the selection list displays the list of drilldowns defined for that resource. The default drilldown is at the top of the list, and the remaining drilldowns are displayed in the sequence as they appear on the Drilldowns tab. This default position is not affected by any sorting of drilldowns.

To change the default drilldown:

1. Open the editor for the data monitor or query viewer you want to edit.
2. On the Drilldowns tab under the Default column, click the button corresponding to the drilldown you want as the default and save.

The default drilldown will appear at the top of the selection list the next time you right-click on the query viewer results or data monitor and select **Drilldown**.

See also ["Sorting or Changing the Order of Drilldowns" below](#) for related information.

Sorting or Changing the Order of Drilldowns

If you create multiple drilldowns to different resource types, the Drilldowns tab displays the drilldowns in the sequence they were created. This initial sort order affects the selection list if you right-click the data monitor or query viewer results on the Viewer panel and select **Drilldowns**.

You can re-order the drilldowns in two ways:

- Sorting the drilldowns
- Moving specific drilldowns up or down the list

To change the sort order:

1. Open the editor for the data monitor or query viewer you want to edit.
2. Click the **Drilldowns** tab and click **Sort** () on the toolbar.

Multiple drilldowns on the Drilldowns tab are sorted in two ways, as follows:

- First, the drilldowns are sorted alphabetically according to resource type: active channels, dashboards, query viewers, and reports.
- Next, within the resource type, drilldowns are again sorted alphabetically by their menu labels.

After you click the **Sort** button, clicking it again will not change the sort order.



Note: Even if the default drilldown moves after sorting on the Drilldowns tab, the default will still be at the top of the selection list when you right-click on the data monitor or query viewer results and select **Drilldowns**. If you want to change the default itself, follow instructions in ["Changing the Default Drilldown" on the previous page](#).

To move a drilldown's position on the list:

1. Open the editor for the data monitor or query viewer you want to edit.
2. Click the **Drilldowns** tab and select a drilldown. Do not click under the Default column if you are not changing the default drilldown.
3. On the toolbar, click the up  or down  arrow buttons to move the drilldown up or down the list.

Removing a Drilldown

You remove any drilldown, including the default drilldown, one at a time. If you delete only the default and you have other drilldowns, the next drilldown on the list becomes the default.

To remove a drilldown:

1. Open the editor for the data monitor or query viewer you want to edit.
2. Click the **Drilldowns** tab.

3. Select the drilldown you want to remove and click **Remove** ( Remove).
4. Repeat as required.

Viewing Query Viewer Results

Where: Navigator > Resources > Query Viewers

To view query viewer results in the default view:

Double-click a query viewer.

The query runs, and returns results in the Viewer on the current state of the network and event flow.

Alternatively, you can add the result of a query viewer directly to a dashboard. For information on this, see ["Adding Query Viewers to Dashboards" on page 372](#).

To view query viewer results in different formats:

1. Right-click the query viewer and select **View Data as** > <Display Format> then choose one of these options:

Display Formats of Query Results

Format	Description
Bar Chart	Display query results as a bar chart.
Horizontal Bar Chart	Display query results as a horizontal bar chart.
Pie Chart	Display query results as a pie chart.
Table	Display query results in table format. Note: Baselines can only be applied to or viewed for query results shown in table format. (For more about establishing and using baselines, see "Defining and Using Baselines" on page 349 .)



Note: By default, chart-style views (Pie and Bar charts) are limited to a maximum of 20 results. Table views can retrieve up to 10,000 rows of data, so it is possible the results in chart views and table views for the same query viewer might not match.


To allow for more results in a chart view, select the **Use classic charts** option in **Global Preferences**. By default, classic charts display a maximum of 99 results. To increase that number, add the following property to the `console.properties` file and specify the desired value:

```
queryviewer.max.dashboard.chart.rows
```

For information about editing the `console.properties` file, see the [ESM Administrator's Guide](#).

Details on how to read and manipulate query results for each of these formats are provided.

- If you select a Table display format, the results are displayed instantly.
- If you select a bar chart or pie chart, you are asked to configure the chart display in the Configure Chart dialog.




Field	Description
Values	The Values drop-down menu lists fields in the query result that contain data types. The value you choose is used as the numbers by which to plot the vertical y axis points on a bar chart or the slice sizes on a pie chart. Values typically represent an unknown set of values, like a count. A common example of numeric data appropriate for values is a time like HourOfDay or a count like Count(Event ID).
Point Labels	The Point Labels drop-down menu provides fields in the query result that contain non-numeric data types. The point labels are used to plot the horizontal x axis labels on a bar chart or the slice labels on a pie chart Examples of non-numeric data types appropriate for point labels are timestamps, strings such as are used for event names, and different types of addresses such as IP or MAC addresses. Point labels are typically a known set of limited values (like hours in a day denoted by timestamps).

2. Select fields for **Values** and **Point Labels**.

Example view settings:


For example, for the Event Counts by Hour of Day query viewer, selecting Count(Event ID) for **Values** (the y axis) and Hour of Day (or Timestamp) for **Point Labels** (the x axis) results in

the following display showing the event count for each hour of the day. The event count is depicted on the vertical y axis, with higher bars representing a higher event count for that hour. The hour of day (time) is represented on the horizontal x axis. The event count is shown for the last 24 hours starting at 11 am.




Understanding the results view:

The results are displayed in the Viewer. The following example shows the “Event Counts By Hour of Day” query result as table, bar chart, pie chart.



Notice that the time range for the base query is shown on the lower left of the query viewer results. Hover the cursor over the time range to see an annotated view of start and end times (data collection start time and data last received). This time range comes from the base query. (Another way to see the query time range is to open the query viewer in the editor and double-click *Query in the Attributes display to drill down to the base query editor, which shows query start and end times.)

Following is the Time Range of Base Query:




Filtering Query Viewer Results

You can filter query viewer results shown in table and chart formats.

To filter query viewer results:

1. Click **"Filter: No Filter"** in the header of the query result view.
Or
Right-click the filter name in the header of the query result view and select **Edit Filter**.
The Common Conditions Editor (CCE) dialog opens.
2. Use the CCE dialog to add a filter. (For details on how to use the CCE dialog to create filters, see the topic on the ["Common Conditions Editor \(CCE\)"](#) on page 887.)



3. Click **OK** to save the filter, and filter the current result view.



Note: Filters on query viewer results are locally saved and available only while the current result set is displayed. These filters are not saved as a part of the query viewer. When you close the query viewer result, the filter is no longer available; recreate it on a new result set.



Tip: You can also apply filters to baseline *delta* columns. See ["Defining and Using Baselines"](#) on page 349.

To remove a filter:

1. Right-click the filter name in the header of the result view
2. Select **Remove Filter**.

Viewing an Event or Resource Directly from the Query Viewer

If your query viewer is for events or resources and the query viewer results include an event ID or resource ID field, you can go directly to a specific event or resource from the query viewer.

Where: Query viewer results

1. Right-click the event or resource.
2. Select **View > [Event Name] Details** or **View [Resource Name] Details**. For example, drill down to
 - An event if the query includes the event ID field
 - An actor if the query includes the actor ID field
 - An asset if the query includes the asset ID field
 - A case if the query includes the case ID field

Working with Query Viewer Results

Various options are available to you with the different query result display formats (Bar Chart, Horizontal Bar Chart, Pie Chart, or Table).

Viewing query results in table format give you the ability to establish baselines and make comparisons, as well as manipulate the table data.



Note: Query viewers and channels display results from variable calculations differently. For example, a value may be displayed as -0.1 in a query viewer, and -0.099999999999... in a channel.

Such variations are due to differences in the way floating point operations are implemented in Java.

Bar charts and pie charts provide at-a-glance, graphical overviews of the results but with fewer options for manipulating the data after the fact.

Other options, such as filtering a query viewer results or running reports, are available on all result views.

Details of working with each view format are provided in the following topics:

- ["Results in Table Format" on the next page](#)
- ["Results in Chart Formats" on page 371](#)

Results in Table Format

To get results in Table format, right-click a query viewer and choose **View Data as > Table**. You can sort, re-order, and create/compare baselines for data in a table view.

Right-click a table column header to get a list of sort and edit options for the column.

Left-click a table column header to sort or reverse-sort it. This affects the entire table.

The arrow next to a column header indicates the column is determining the current sort order for the table and shows the current sort order. (For example, this column is sorted showing highest count first.)

You can add a baseline to a **table** view of query result data, then run **baseline** comparisons to identify deltas in network behavior.

Only **table views** of query result data can be baselined and compared.

Timestamp	Name	Count(Event ID)	Count(Event ID) - (Reference Baseline)
16 Apr 2008 13:00	Sort Column	5460	5375
16 Apr 2008 13:00	Monitor Event	5990	772
16 Apr 2008 13:00	ip value count data monitor value...	773	5819
16 Apr 2008 13:00	Show Column	5811	42
16 Apr 2008 13:00	Size Column to Fit	5704	33
16 Apr 2008 13:00	Help	5685	3
16 Apr 2008 13:00:00 PDT	Count(Event ID)	5630	0
16 Apr 2008 13:00:00 PDT	Count(Event ID) - (Typical Weekday Baseline)	5460	0
16 Apr 2008 13:00:00 PDT	Count(Event ID) - (Reference Baseline)	5460	0
16 Apr 2008 13:00:00 PDT	Starting Trend Query	96	0
16 Apr 2008 13:00:00 PDT	Sidetable Sizes (Rows)	30	0

See also:

- [""Analyze in Channel" Options on the Table View" below](#)
- ["Column Sort, Display, and Edit Options" on page 370](#)

"Analyze in Channel" Options on the Table View

The following right-click **Analyze in Channel** options are available on query viewer results in **table** format (obtained by choosing **View Data as > Table**):

- **Baselines.** Right-click *anywhere on the table of results* in the Viewer to add a baseline or compare the current results to an existing baseline.
- **Drilldowns.** Right-click a *row* in the table result to launch a given drilldown on that row item (if drilldowns are provided in the query viewer).
- **Channels.** Right-click a *cell* in the table result to create an active channel with a filter based on the value of the selected table cell.
- **Conditions.** Right-click a *cell* in the table result to add a filter condition based on the value of the cell.

 **Note:** The **Analyze in Channel** option is not available if your base query is not on events; for example, if the query is on a session list.


Analyze in Channel Options for Results are shown in the following table.

Analyze in Channel Options on Query Viewer Results

Option	Description
Add as baseline	Add the current results as a baseline for the query viewer. Right-click <i>anywhere on the table result</i> in the Viewer to add a baseline to the query viewer or compare the current results to an existing baseline. (See "Defining and Using Baselines" on page 349 and "Adding a Baseline" on page 351.)
Compare with: <Baseline>	Compare the current results with the selected baseline. Right-click <i>anywhere on the table result</i> in the Viewer to compare the current results to an existing baseline. This menu option is available if there is one or more baselines established for the query viewer. All baselines associated with the query viewer are available from this menu for comparison. (See "Comparing Displayed Results to a Baseline" on page 352.)
Drilldowns	See "Managing Drilldowns from Query Viewers" on page 356 and "Drilldown Example" on page 384. Query viewers enable you to drill down to Active Channels, dashboards, query viewers, and reports. If the query includes an event ID or resource ID, you can also drill down to that resource. See "Viewing an Event or Resource Directly from the Query Viewer" on page 367 for details. If there are drilldowns associated with the query viewer, these are listed after the baseline options on the right-click Analyze in Channel menu for a selected row in the query viewer result. Right-click a row in the table result, and choose Analyze in Channel > <Drilldown Option>. For example, an Events query viewer could provide drilldowns to view all source addresses for a selected event. Assuming each row in the result table represents an event, choosing this drilldown from the Analyze in Channel menu would lead to a table showing source addresses for the selected event.
Create Channel	Creates an active channel with a filter based on the selected cell in the table result. For example, right-clicking a table cell with an event name and choosing Analyze in Channel> Create Channel [EventName] creates an active channel that monitors and filters for occurrences of that event name. The filter is always set to the value of the cell (which in this example would be the event name). For more information about creating active channels, see "Creating or Editing an Active Channel" on page 224.
Add Condition	Brings up the Conditions Editor for the selected item, where you can add or modify conditions (filters) on the selected item. Right-click a <i>cell</i> in the query viewer table result to add a filter condition based on the value of the cell. For more information on working with Conditions, see "Common Conditions Editor (CCE)" on page 887.

Column Sort, Display, and Edit Options

Right-click a column header in a query viewer table result to get various options on that column.



Column Header Options

Option	Description
Sort Column	Sort items in the column in ascending or descending order. Columns that have been sorted after the query viewer run show an up or down arrow next to them to indicate the direction of the sort. You can also sort the column by left-clicking the column header. Clicking multiple times toggles the sort between: ascending order (indicated by a up arrow next to the header) descending order (indicated by a down arrow next to the header) Notes: - Sorting on the contents of a column after a query viewer displays its results changes the view of the data provided by the original query. A query sorts during a query run, and then displays the data based on the sorting it did. If you click columns to re-sort, you are changing the sort order the query gave you. In the cases where the original query used a “single-column” sort, you can “get back” to it in the viewer, but you can’t get back to a multi-column sort because this is offered only in the query sort options, not on the Console UI. - Keep in mind that this option sorts on the data result returned by the query. This in combination with query row limits (applied when the query is run) can sometimes yield unexpected results. Example: If the query is defined to run on 2 days’ worth of data but hits the 10,000 row limit after processing only 1 day of data, then only 1 day’s worth of data is returned in the result. An “after-query” sort, in this example, is a sort on only 1 day’s worth of data. - Sorting at the query viewer level sorts only the data returned by the query to Viewer. Initial sorting is done by the base query, which is responsible for running against the database. If the query level sort is yielding unexpected results, keep in mind that the original base query sort determines how much you can modify the view of the result. See also "Sort Baseline Data" on page 354.
Remove Sort	Remove a sort on the selected column. You can remove sorting imposed when the query viewer was run or when a UI column-click sort was done on the displayed result.

Column Header Options, continued

Option	Description																																				
Show Column	Right-click anywhere on any column header in a table to get a context menu of columns included in the display result. Select columns to hide or show in the result. Columns with no checkmark beside them are hidden. This is the equivalent of hiding or showing a column before the query viewer runs. (However, only columns configured to be included in the original query are available to hide/show after the query is run.) To show a column in the results view that is currently hidden (whether before or after the query ran), right-click again and choose it (checkmark it). See also "Show or Hide Baseline Columns" on page 353.																																				
Size to Fit	Expand the column, if needed, to accommodate the full width for text in each row of the selected column.																																				
Drag-and-Drop options	Left-click-and-drag on a column header to reposition it in a different horizontal order in the table. For example, if the original query viewer result shows columns in this order: <table><tr><th>TimeStamp</th><th>Name</th><th>Count(Event ID)</th></tr><tr><td>12 May 2008 18:00...</td><td>Monitor Event</td><td>5460</td></tr><tr><td>12 May 2008 18:00...</td><td>Top value count data monitor value current</td><td>807</td></tr><tr><td>12 May 2008 18:00...</td><td>NETBIOS SMB-D5 DCERPC NTLMSSP asn1 ...</td><td>661</td></tr><tr><td>12 May 2008 18:00...</td><td>NETBIOS SMB-D5 Session Setup AndX req...</td><td>658</td></tr><tr><td>12 May 2008 18:00...</td><td>Task successfully scheduled</td><td>40</td></tr></table> You can click-and-drag TimeStamp to the right so that the columns display in this order: <table><tr><th>Name</th><th>TimeStamp</th><th>Count(Event ID)</th></tr><tr><td>Monitor Event</td><td>12 May 2008 18:00...</td><td>5460</td></tr><tr><td>Top value count data monitor value current</td><td>12 May 2008 18:00...</td><td>807</td></tr><tr><td>NETBIOS SMB-D5 DCERPC NTLMSSP asn1 ...</td><td>12 May 2008 18:00...</td><td>661</td></tr><tr><td>NETBIOS SMB-D5 Session Setup AndX req...</td><td>12 May 2008 18:00...</td><td>658</td></tr><tr><td>Task successfully scheduled</td><td>12 May 2008 18:00...</td><td>40</td></tr></table>	TimeStamp	Name	Count(Event ID)	12 May 2008 18:00...	Monitor Event	5460	12 May 2008 18:00...	Top value count data monitor value current	807	12 May 2008 18:00...	NETBIOS SMB-D5 DCERPC NTLMSSP asn1 ...	661	12 May 2008 18:00...	NETBIOS SMB-D5 Session Setup AndX req...	658	12 May 2008 18:00...	Task successfully scheduled	40	Name	TimeStamp	Count(Event ID)	Monitor Event	12 May 2008 18:00...	5460	Top value count data monitor value current	12 May 2008 18:00...	807	NETBIOS SMB-D5 DCERPC NTLMSSP asn1 ...	12 May 2008 18:00...	661	NETBIOS SMB-D5 Session Setup AndX req...	12 May 2008 18:00...	658	Task successfully scheduled	12 May 2008 18:00...	40
TimeStamp	Name	Count(Event ID)																																			
12 May 2008 18:00...	Monitor Event	5460																																			
12 May 2008 18:00...	Top value count data monitor value current	807																																			
12 May 2008 18:00...	NETBIOS SMB-D5 DCERPC NTLMSSP asn1 ...	661																																			
12 May 2008 18:00...	NETBIOS SMB-D5 Session Setup AndX req...	658																																			
12 May 2008 18:00...	Task successfully scheduled	40																																			
Name	TimeStamp	Count(Event ID)																																			
Monitor Event	12 May 2008 18:00...	5460																																			
Top value count data monitor value current	12 May 2008 18:00...	807																																			
NETBIOS SMB-D5 DCERPC NTLMSSP asn1 ...	12 May 2008 18:00...	661																																			
NETBIOS SMB-D5 Session Setup AndX req...	12 May 2008 18:00...	658																																			
Task successfully scheduled	12 May 2008 18:00...	40																																			

Results in Chart Formats

To get results in Chart format, right-click a query viewer and choose either:

- **View Data as > Bar Chart or Horizontal Bar Chart**
- **View Data as > Pie Chart.**

Chart Options

Option	Description
Drilldowns	Query viewers can provide <i>drilldowns</i> to Active Channels. If the query includes an event ID or resource ID, you can also drill down to that resource. See "Viewing an Event or Resource Directly from the Query Viewer" on page 367 for details. If there are drilldowns associated with the query viewer, select an item in the first or key column, then right-click to get drilldown options in the Analyze in Channel menu. For example, an Events query viewer could provide drill-downs to view all source addresses for a selected event. Choosing this drilldown from the Analyze in Channel menu on a query result would lead to a table showing source addresses for the selected event. See "Managing Drilldowns from Query Viewers" on page 356 and "Drilldown Example" on page 384.
Create Channel	Creates a channel on the selected item. (For example, right-clicking an event and choosing Analyze in Channel > Create Channel [EventName] creates an active channel that monitors and filters for occurrences of that event. For more information about creating active channels, see "Creating or Editing an Active Channel" on page 224.
Add Condition	Brings up the Conditions Editor for the selected item, where you can add or modify conditions (filters) on the selected item. For more information on working with Conditions, see "Common Conditions Editor (CCE)" on page 887.

Troubleshooting Query Viewers

If queries time out, especially if your environment monitors high event rates (in the thousands per second), try reducing the number of rows to the range of 100 to 1000 and see if there is an improvement. If that does not improve execution time, refer to the *ESM Administrator's Guide's* Troubleshooting chapter. Read the topic, *Query and Trend Performance Tuning*.

Adding Query Viewers to Dashboards

You can add a query viewer result to a dashboard as follows:

Where: Navigator > Resources > Query Viewers

1. If you have identified an existing dashboard to which you want to add the query viewer, open the dashboard in the viewer and make sure it is the focus. If you want to add the query viewer to a new dashboard, continue to the next step.
2. Right-click a query viewer and select **Add to Dashboard As >**, then select an applicable display format (see ["Dashboard Display Formats" on page 259](#)).

The query viewer result is displayed on the open dashboard. If a dashboard is not displayed, a new untitled dashboard is created for the query viewer result.

3. Save the existing dashboard.

Or if this is a new dashboard:

- a. Right-click the title bar of the dashboard and choose **Save Dashboard As**.
- b. In the popup dialog, navigate to the group where you want to save the dashboard, enter a name for the dashboard, and click **OK**.

By default, this new dashboard is a regular dashboard. If you want to change it to a custom view dashboard, see ["Using Custom View Dashboards" on page 261](#).

You can add multiple query viewer results sets along with other resources to a single dashboard.

For more information about working with dashboards, see ["Using Dashboards" on page 252](#).

Adding Query Viewers as Startup Views

Where: Navigator > Resources > Users

To set up Query Viewers as the startup view for a group:

1. Right-click a user group and select **Edit Group**.
2. In the editor for the selected group, click **Startup Views** tab, then click the **Query Viewers** subtab.

3. Click **Add** ( **Add...**).

4. In the Query Viewer Selector , navigate to and select (checkmark) the query viewer you want as the startup query viewer for this group. Then click **OK**

The full path to the query viewer you selected is shown on the Query Viewers tab in Startup Views.

5. Click **OK**.

For more information on editing groups and startup views, see ["Managing User Groups" on page 99](#).



Tip: Regardless of startup view settings for groups, if Query Viewers are showing when you quit the Console, these are reloaded when you restart the Console.

Generating Reports from Query Viewers

After you run a query viewer, you can generate a simple report containing the results. Reports initiated from query viewers are provided for convenience as a quick way to share the result data. Query viewer reports are limited to displaying data from the single query covered by the query viewer and retain the format of the chart or table in which the query viewer results are displayed. For information on creating and publishing richer, highly formatted reports on multiple data sources see ["Building Reports" on page 388](#) and ["Running and Managing Reports" on page 446](#).



Tip: On row limits


The report display format is based on the display chosen for the query viewer result. For example, if you chose to view query data as a *pie chart*, the generated report shows the same pie chart view. To generate a report showing results for the same query as a *bar chart* or *table*, you must then re-run the query viewer (<Query Viewer> > View Data as) in one of those formats, and then generate the report from that view.

The report contents might not include as much data as the query viewer result shown in the Console for these reasons:

- Reports on pie charts and bar charts have a default row limit of 25 up to 99 rows as the maximum. This is user-configurable. You can set a higher or lower row limit on the Report Parameters dialog you get when you run the report. (See the procedure below.)
- Reports on *tables* can accommodate up to 10,000 rows.

To generate a report on a query viewer:

1. Right-click the query viewer results table or chart (anywhere in the Viewer panel) and click **Report**.



2. Specify the options on the Report Parameters dialog or take the defaults and click **OK**. For more help on setting report parameters, see ["Setting Default and Custom Report Parameters" on page 415](#).



Tip: If you click **Save Output** on the Report Parameters dialog, you get additional options for setting the archived report under the Save Output Parameters section.

3. When the report is ready, a dialog gives you the option of opening it to view it now or saving it to a location you specify through a file browser.

Select **Open** to view the report or **Save** to save it in a specified location.

Example Queries for Common Scenarios

Query viewers can be used to monitor daily network traffic and get high level summaries of typical activity. Query viewers can also be used to drill down on anomalies or other interesting events.

Following is a brief, conceptual scenario of how an **analyst** might use query viewers to **monitor and investigate** certain types of activity.

Also included here is a description of how the **query content developer** might **build and configure** the base query and query viewers that the analyst uses.



Tip: In practice, ArcSight ships with pre-built queries and query viewers as standard content. It is likely that the types of resources described here are provided with ArcSight.

Even so, the configuration of the base query and query viewers is described to illustrate and support this example, and show how a content developer might fine tune these resources to gather the information needed.

Basic Analysis High Level Summaries

A security analyst wants to check if anything unusual is happening on their system. The analyst brings up a query viewer called “Events” that shows all events by event name for the last 2 hours. The columns include:

- Event name
- Total count of all events
- Count by unique source address
- Count by unique destination address

Analyst’s First View of Events

The analyst can easily glance at the data and see if anything looks out of the ordinary. Columns can be sorted and filters can be changed to refine the details. The data should come up almost immediately.



The screenshot shows the 'Viewer' window with the 'Events: Table' tab selected. The query is 'Events' and it shows 114 matches. The last update is '18 Jun 2017 10:08:20 PDT'. The filter is 'No Filter'. The table has four columns: 'Name', 'COUNT(Events)', 'COUNT(Source ...)', and 'COUNT(Destina...)'.

Name	COUNT(Events)	COUNT(Source ...)	COUNT(Destina...)
ActiveList entry expired	338	0	1
ActiveList entry updated	2979	0	1
AddToList: Success	3213	0	1
Agent [Nifty Event Player] reconnected	1	1	1
Agent Login	2	1	1
Agent [Nifty Event Player] heartbeat time...	1	0	1
Agent [Nifty Event Player] type [testalert...	1	0	0
Agent updated	2	1	1
Application Event Counts	8	0	0
ArcSight Event Flow Statistics	9	0	0
ArcSight Manager Started	1	0	1
ArcSight User Login	1	1	0
Asset updated	40	0	1
AttachmentOnlyGroup [Package] inserted	1	1	1
Attack From Suspicious Source	2	1	1
Attack Rates by Attacker Zone	2	0	0
Attack Rates by Service	2	0	0
Attack Rates by Targeted Zone	2	0	0
Attacker Zones by Service	2	0	0
Case inserted	2	0	1
Case updated	2	0	1
Channel [Last Hour] got attached	1	1	1
Channel [Last Hour] query completed	2	0	1
Channel [Live] got attached	1	1	1
Channel [Live] query completed	2	0	1
Compromise - Success	2	1	1
Connector Down	1	0	0
Connector Still Down	2	0	0
Connector Up	1	0	0
CreateNewCase: Success	2	0	1
Data Monitor Entry Timeout	1	0	1


How the Events Query Viewer is Built

The **Events query viewer** described in this example leverages the Events query.

Attributes


Bringing up the *query viewer editor* for the **Events query viewer** shows that the Events query is used as the base query. Bringing up the **Events query** (base query) in the *query editor* shows that the base query searches on events for the last 2 hours. (Queries are under Reports > Queries in the Navigator.)

The Events query viewer uses the Events query as its *base query*.




Fields

The fields selection, order by, and group by logic are all defined in the Fields tab for the base query. The Events query viewer inherits the fields from the base query. These show up on the query viewer Fields tab.



Events Base Query Conditions Tab

The condition logic to search on Events is defined in the Conditions tab for the base query.




Note: If the event value in your query is the @ symbol by itself, enclose it in double quotes. For example:

Name Contains "@"

You are not required to use the double quotes if the @ symbol is used with other text, for example, **Name Contains @mycompany**.

Analyst's First View of Events

The analyst can easily glance at the data and see if anything looks out of the ordinary. Columns can be sorted and filters can be changed to refine the details. The data should come up almost immediately.



The screenshot shows the 'Viewer' window with the 'Events: Table' tab selected. The query is 'Events' and it shows 114 matches. The last update is '18 Jun 2017 10:08:20 PDT'. The filter is 'No Filter'. The table has four columns: 'Name', 'COUNT(Events)', 'COUNT(Source ...)', and 'COUNT(Destina...)'.

Name	COUNT(Events)	COUNT(Source ...)	COUNT(Destina...)
ActiveList entry expired	338	0	1
ActiveList entry updated	2979	0	1
AddToList: Success	3213	0	1
Agent [Nifty Event Player] reconnected	1	1	1
Agent Login	2	1	1
Agent [Nifty Event Player] heartbeat time...	1	0	1
Agent [Nifty Event Player] type [testalert...	1	0	0
Agent updated	2	1	1
Application Event Counts	8	0	0
ArcSight Event Flow Statistics	9	0	0
ArcSight Manager Started	1	0	1
ArcSight User Login	1	1	0
Asset updated	40	0	1
AttachmentOnlyGroup [Package] inserted	1	1	1
Attack From Suspicious Source	2	1	1
Attack Rates by Attacker Zone	2	0	0
Attack Rates by Service	2	0	0
Attack Rates by Targeted Zone	2	0	0
Attacker Zones by Service	2	0	0
Case inserted	2	0	1
Case updated	2	0	1
Channel [Last Hour] got attached	1	1	1
Channel [Last Hour] query completed	2	0	1
Channel [Live] got attached	1	1	1
Channel [Live] query completed	2	0	1
Compromise - Success	2	1	1
Connector Down	1	0	0
Connector Still Down	2	0	0
Connector Up	1	0	0
CreateNewCase: Success	2	0	1
Data Monitor Entry Timeout	1	0	1


How the Events Query Viewer is Built

The **Events query viewer** described in this example leverages the Events query.

Attributes


Bringing up the *query viewer editor* for the **Events query viewer** shows that the Events query is used as the base query. Bringing up the **Events query** (base query) in the *query editor* shows that the base query searches on events for the last 2 hours. (Queries are under Reports > Queries in the Navigator.)

The Events query viewer uses the Events query as its *base query*.




Fields

The fields selection, order by, and group by logic are all defined in the Fields tab for the base query. The Events query viewer inherits the fields from the base query. These show up on the query viewer Fields tab.



Events Base Query Conditions Tab

The condition logic to search on Events is defined in the Conditions tab for the base query.



Note: If the event value in your query is the @ symbol by itself, enclose it in double quotes. For example:

Name Contains "@"

You are not required to use the double quotes if the @ symbol is used with other text, for example, **Name Contains @mycompany**.


How the Events Query Viewer is Built

The **Events query viewer** described in this example leverages the Events query.

Attributes


Bringing up the *query viewer editor* for the **Events query viewer** shows that the Events query is used as the base query. Bringing up the **Events query** (base query) in the *query editor* shows that the base query searches on events for the last 2 hours. (Queries are under Reports > Queries in the Navigator.)

The Events query viewer uses the Events query as its *base query*.




Fields

The fields selection, order by, and group by logic are all defined in the Fields tab for the base query. The Events query viewer inherits the fields from the base query. These show up on the query viewer Fields tab.



Events Base Query Conditions Tab

The condition logic to search on Events is defined in the Conditions tab for the base query.



Note: If the event value in your query is the @ symbol by itself, enclose it in double quotes. For example:

Name Contains "@"


You are not required to use the double quotes if the @ symbol is used with other text, for example, **Name Contains @mycompany**.

Drilldown Example

Continuing from the example in ["Basic Analysis High Level Summaries" on page 376](#), the security analyst notes that one of the counts seems troublesome. For example, "Attack from Suspicious Source" is high and showing a lot of unique destination addresses. The analyst would right-click this row and choose **Show Source Addresses**.

The resulting query viewer would show, for this event and time range, the source addresses, as well as other columns of interest (for example, destination address). Then by sorting by source address, the analyst could decide if a single source address (probably with the highest count)

was the initiator of most of the attacks. This information could also be provided from an appropriate back end trend table (the same one or a different one), and, as a result, the display should come up almost immediately.




The screenshot shows the 'Viewer' window with the title bar 'Viewer'. Below the title bar is a tab labeled 'Drilldown on Sources: Table'. The main content area displays the following information:

- Query: Drilldown-src
- Last Update: 1 Aug 2008 18:24:12 PDT
- Filter: Name = "ICMP PING"
- 13 shown / 13 matches

Source Address	COUNT(Source Address)
192.168.1.1	1
192.168.1.2	482
192.168.1.3	1
192.168.1.4	1
192.168.1.5	2
192.168.1.6	1
192.168.1.7	1
192.168.1.8	1
192.168.1.9	1
192.168.1.10	1
192.168.1.11	1
192.168.1.12	1
192.168.1.13	1

The status bar at the bottom shows the time range '8/1 16:22:58 - 8/1 18:22:58'.

The analyst could also show **destination** addresses for the same event row, if that drilldown is defined as a part of the query viewer.



The screenshot shows the 'Viewer' window with the title bar 'Viewer'. Below the title bar is a tab labeled 'Drilldown on Destinations: Table'. The main content area displays the following information:

- Query: Drilldown-dst
- Last Update: 1 Aug 2015 18:31:53 PDT
- Filter: Name = "ICMP PING"
- 7 shown / 7 matches

Destination Address	COUNT(Destination Address)
192.168.1.1	1
192.168.1.2	4
192.168.1.3	3
192.168.1.4	2
192.168.1.5	2
192.168.1.6	1
192.168.1.7	577

The status bar at the bottom shows the time range '8/1 16:31:43 - 8/1 18:31:43'.

How the Console Builds Drilldowns

The source and destination drilldowns are added to the Events query viewer on the Drilldowns tab at content development time.

Here is the Drilldowns > Field Mapping tab for the Events query viewer example. The drilldown requires that at least one field is mapped. In this example, this is the Name field.

Field Mapping		
Display Sort		
+ - x		
Source Field	Condition	Destination Field
Name	=	Name

Here is the Drilldowns > Display tab for the Events query viewer example.

Field Mapping		
Display Sort		
+ - x		
☒	Name	Display Name
☒	Name	Name
☒	Source Address	Source Address
☒	COUNT(Source Address)	COUNT(Source Address)

Non-Event Analysis Example

A security analyst wants to examine “Asset Counts by Vulnerability.” The analyst selects this viewer and gets the most recent result (from a trend run) and can examine a table containing columns: Vulnerability and Asset Count. Right-clicking a particular vulnerability row would allow drilldown into the assets with that vulnerability.

Baseline Analysis for Data Comparison

The security analyst notes that one of the counts seems significantly higher than last recalled. The analyst right-clicks the query viewer and selects “Compare with Baseline”, from which there are zero or more baselines to choose.

This makes additional columns available to the currently displayed viewer that can be added by the user. For example, a new column could be added next to the current “Count” column showing “Count - <Selected Baseline>”. This is a comparison number showing the difference between the current value of the count and the baseline value for the count. This is positive, negative, or empty (if a baseline doesn't exist for this vulnerability). The analyst can right-click the new column to sort this column in ascending or descending order.

Other options available to the analyst would be:

- **Add as Baseline...** to save the current values in the display as the new named baseline.
- **Compare with...** to compare to any other set of data available in the trend table.

History Analysis Example

As hinted in this example, any previous trend runs can be used for baseline comparison. Similarly, the analyst can change the query viewer to go back into the past to look at previous data. The analyst could use the default baseline and go back in history to see when some count began to significantly differ from the baseline.


Chapter 13: Building Reports

Reports are captured views or summaries of data that can be viewed in the ArcSight Console or exported for sharing in a variety of file formats. Reporting is an essential tool for communicating the state of your enterprise security to internal and external stakeholders.

Reporting is a broad subject in ArcSight. Because it can use all the scheduling, conditional logic, resource- and rules-based filtering capabilities of the system, the possibilities can take some time to explore. Creating a report is a multi-step process that can involve steps using several different resources.

When you have source data defined in queries or trends, you can create reports to present the data in charts and tables. Use one of the templates provided with ArcSight.

The reports resource defines how query data is bound to a report template. Depending on the report template you use, the reports editor exposes different parameters, variables, and conditions that enable you to choose which elements of the query data you want to show in the report. You can also apply additional functions to run on the data, and set numerous formatting options.



See also ["Running and Managing Reports" on page 446](#) and ["Archiving and Scheduling Reports" on page 459](#).

Understanding the Reporting Workflow


Building reports is a multi-step process that involves use of a few different data gathering and reporting tools. ArcSight can gather report data using standard queries or trends.



Tip: Reports can be relatively simple (you can create a report with the Report Wizard based on the results of a single query) or complex (you can create a report based on the results of layers of queries and trends that feed data results up the chain as the basis for new queries). See ["End-to-End Reporting Examples" on page 430](#) for examples of both basic and complex reports.

Following is a quick overview of the reporting workflow tasks and tools, along with a reminder about dependencies among reporting resources.

For a more in-depth description of how these elements build on each other to create various views of the data, see also ["Query-Trend Relationships in Reporting" on page 468](#).



Step 1 - Build a Query

A query is an ArcSight resource that defines the parameters of data to gather from an ArcSight data source. The results of the query then become the basis for one or more ArcSight reports or trends. As a data source, queries can use the database of events, assets, cases, notifications, active lists, session lists, or data gathered from a trend.

Queries are described in ["Queries" on page 319](#).



Note: If all you want to do is build a report based on a single query, at this point you can skip to step 4 and select a template. (See ["Step 4 - Select or Design a Report Template" on the next page](#).)



Tip: Queries built for reports can also be used in query viewers.

Query viewers provide several advantages. If you want to run quick SQL queries for monitoring and analysis outside of the reporting resource, you can use query viewers. You can add query viewers to dashboards and generate simple reports on query viewer results.

For information on query viewers, see ["Query Viewers" on page 340](#).

Step 2 - Build a Trend Based on a Query

A trend is an ArcSight resource that defines how and over what time period data is evaluated for trends. A trend is always based on a query, and this query is called the base query for the trend. The trend results are stored in a trend table in the database, and these trend tables can be queried. Queries on trend tables are called trend queries. Trends can also be used as the primary data source for a report.

Trends are described in ["Building Trends" on page 467](#).



Note: If you want a report based on a single trend, at this point you can skip to step 4 and select a template. See ["Step 4 - Select or Design a Report Template" below](#).

Step 3 - Build a Query Based on a Trend

At this point, you have the option of using a simple query or trend in a report, or you can further refine query results by using a trend in another query (a trend query).

See the ["Building a Query" on page 320](#) and ["Building Trends" on page 467](#) for more information on how to do this.

Data Gathering: Query Building	Data Gathering: Trend Defining
 Query source data - Specify the data you want to work with - Narrow the results by setting conditions and variables	 Is data a trend? - Design interval trend to operate on events - Design snapshot trend to operate on assets, network model, cases, and notifications

Step 4 - Select or Design a Report Template

Use an existing report template layout or create your own using the new Report Designer tool. For information on working with templates, see ["Using Report Templates" on page 426](#).

Layout Designing
 Design report template - Using a stock template? - Design your own: chart, table, combination

Step 5 - Create a Report

A report is an ArcSight resource that binds data from a query or trend to an existing report template. Once run, the results of a report can be viewed in the ArcSight Console's Viewer panel, saved (archived), and/or exported in a variety of formats. Reports can be scheduled to run at regular intervals, and also can be run on demand as needed.

Reports are described in ["Creating or Editing a Report" on page 393](#), and an overview of the whole topic is provided in ["Understanding the Reporting Workflow" on page 388](#).

Focused reports enable you to run the same report definition on different subdivisions of the data without having to copy and modify the master report every time. For example, you can run an individual Top 10 Infected Systems report for each of your business divisions.

The job scheduler enables you to schedule reports and focused reports to run automatically at specific time intervals. (The job scheduler is also used as a part of building trends which, by nature, involve scheduling.)

Report Building	Report Refining
 Bind the query data to the report template - Refine or customize the query by choosing specific elements of the data - Add functions - Format, organize, and re-order results - Format fonts, text color, and background color	 Create focused reports - Optional - Narrow the focus of a report using focused reports - Run the same report with different parameters for each division



Note: Queries and trends are intended to capture data. Reports are used to display the data from queries and trends. For example, if you wanted to run monthly or quarterly reports on VPN login statistics, you would first create one or more queries to capture the data, then create trends (based on the queries) to define a schedule for running the queries and storing the results, and finally create and run reports on the trends. For a full walk-through of this process, see ["End-to-End Reporting Examples" on page 430](#).

Step 6 - Run a Report

ArcSight ships with a set of ready-made reports available under the Reports resource. For example, on the Navigator panel under the Reports resource look in /All Reports/ArcSight Administration/. Open the sub-groups (folders) to see provided reports.

For information on how to run an existing report, see ["Running and Managing Reports" on page 446](#) and Running a New or Archived Report.

Step 7 - Archive and Maintain Reports

After running a report, you have the option to save (archive) the report results. This enables you to retrieve a particular report for immediate viewing without having to regenerate the report. Reports that are run on demand are saved on the Archives tab just like scheduled reports. If the Save Output option is chosen for an on-demand report, the archived report has an expiration date of 6 months from the time it was run (by default). If the Save Output option is not chosen for an on-demand report, the report is maintained in the archive for one day only.

Archived reports can also be sent to a notification group after the scheduled report is run.

For information on how to archive and maintain reports, see ["Archiving and Scheduling Reports" on page 459](#).



Managing Dependencies for Reports Resources

Queries, trends, and reports generally have multiple dependencies upon each other. Modifying some elements within one resource can affect another. If modifications to a resource impact another to the extent that the dependent resource is rendered unusable, errors are reflected in the ArcSight Console. ESM manages and updates most of these resources and dependencies automatically, but not all.

For example, a trend built on a query relies on a set of fields (columns) contained in the base query. If you modify fields in the base query that are used in the trend, the trend is disabled. Similarly removing a resource (like a query) that another resource (like a report) depends on generates error messages on the ArcSight Console.



Caution: Instead of modifying the trend's base query, the proper approach is to create a different query, then create a new trend using that query.

Creating or Editing a Report

**Caution:** Do not exceed more than 10,000 resources in a group.

Where: Navigator > Resources > Reports > Reports tab

The high-level steps for creating a report are as follows:

1. To create a report, right-click a reports group and select **New Report**.

**Note:** As a general rule, create content in the user's own folder.

To edit a report, right-click the report and select **Edit Report**.

2. In the edit panel, define report attributes such as report name (required), an optional alias, and other details. See ["Defining Report Attributes" on the next page](#).

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

3. Select the report template you want to use. Defaults are provided. See ["Using Report Templates" on page 426](#).
4. Select report data by:
 - a. Associating one of these existing data sources: a query, a trend, a session list, or an active list. You choose from the displayed resource selector.
 - b. Specifying what parts of the query data source you want to use for each report element. Optionally, apply legends and top/bottom functions.

See ["Binding Data to the Report" on page 398](#).

Note: Unlike report templates and parameters which provide defaults, you cannot save the report without specifying the data source.

5. Specify report parameters concerning output details, such as file format, paper size, and routing instructions. You can set limits on the query return, such as row limits, time zone restraints, apply filters, and specify report start and end times. See ["Setting Default and Custom Report Parameters" on page 415](#).
6. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
7. Click **OK**.



Note: Be sure to click **Apply** or **OK** frequently to save settings intermittently as you work through the above steps. Clicking **Apply** saves settings and leaves the Editor open. Clicking **OK** saves settings and closes the Editor for this query. If you do not apply or accept settings using these buttons, your settings are not saved.

8. Run the report to test it as described in ["Running a Report" on page 446](#).

The next sections provide details on how to use the Report editor to define report attributes, apply a template, choose report data, and specify report parameters.

Defining Report Attributes

Report settings include general attributes, template selection, and data specification.

The **Report Attributes** tab is where you define a report name, set alias report name and notification options, and view tracking details such as when the report was created and last updated.

The following field is required on the Report Editor's **Attributes** tab.

Report Field	Description
Name	Default name for the report. Spaces and special characters are allowed. Note: You can change the report name in the following ways: • Edit the report and changing the name on the Attributes tab. • Set the report name on the Jobs tab where you define the report's archiving schedule (see "Archiving and Scheduling Reports" on page 459). • Change the name at report runtime (see "Running a Report" on page 446).



Tip: Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

The following example shows the completed Report Attributes tab for a report:

The screenshot shows the 'Inspect/Edit' dialog box for a report titled 'Report:VPN Login Outcome Trend...'. The 'Attributes' tab is selected, displaying the following fields:

Report	
Name	VPN Login Outcome Trend Report
Common	
Resource ID	9B10-PIABABCgWaltZ021bw==
External ID	
Alias (Display Name)	
Description	ed on the trend query on successful logins.
Version ID	
Deprecated	☐
Assign	
Owner	jAnalyst
Notification Groups	SOC Operators
Parent Groups	
jAnalyst's Reports	/All Reports/Personal/jAnalyst's Reports/
Creation Information	
Created By	admin
Creation Time	6 Oct 2015 10:42:19 PDT
Time Since Creation	10 min(s) 1 sec(s)
Last Update Information	
Last Updated By	admin
Last Update Time	6 Oct 2015 10:51:28 PDT
Time Since Last Update	52 sec(s)

At the bottom of the dialog, there is a section for '(Name)' and '(Description)', and buttons for 'Preview...', 'OK', 'Cancel', 'Apply', and 'Help'.


Report Templates

The **Template** tab is where you specify the template for the report. You can specify fonts, colors, page headers and footers, and the chart and table combinations and layout you want to use.

Report Template Selection

To populate the editor, select a template from the **Report Template** drop-down menu. ArcSight comes with six stock templates in the System templates folder, or you can navigate to your own template.

The following example shows the system template Three Charts Landscape.



Report Template:
Choose a template
(ArcSight-provided or
custom template)

Text Components:
Specify Header,
Footer, Title, Subtitle,
and so on; and text
formatting options for
the report.

Preview area:
Shows the layout of
the Report Template
you chose.

Text Attributes

The **Text Attributes** section provide fields to specify values for each set of text attributes.

Report Template Text Components

Attribute	Description
Header	
Text	Enter the text you want to use as the header of your the pages in your report, such as the name of your department, or the series of reports to which it belongs. Note: You can use Velocity template references for fields that accept text, as described in " Velocity References for Reports " on page 1091.
Horizontal Alignment	From the drop-down menu, select where you want the header to appear in the header area: left, right or center.
Vertical Alignment	From the drop-down menu, select where you want the header to appear in the header area: top, center, or bottom.

Report Template Text Components, continued


Attribute	Description
Font	From the drop-down dialog, select a font from the list of fonts available on your local system, font size, and style (bold, italic). The preview window indicates how the font will look.
Foreground Color	From the drop-down dialog, select a foreground color. This is the color of the lettering.
Background Color	From the drop-down dialog, select a background color. This color fills the header box.
Footer: Footer, Disclaimer	
Text	Enter the text you want to use as the footer of your the pages in your report, such as the name of your company, a confidentiality statement, or the date. You can use the variables provided (such as \$currentpagenumber and \$totalpagenumber for page numbers). These are evaluated when you run the report to populate report output with appropriate numbering. Note: You can use Velocity template references for fields that accept text, as described in "Velocity References for Reports" on page 1091. Enter text for Footer: Disclaimer as required by your company.
Horizontal Alignment	From the drop-down menu, select where you want the footer to appear in the footer area: left, right or center.
Vertical Alignment	From the drop-down menu, select where you want the footer to appear in the footer area: top, center, or bottom.
Font	From the drop-down dialog, select a font from the list of fonts available on your local system, font size, and style (bold, italic). The preview window indicates how the font will look.
Foreground Color	From the drop-down dialog, select a foreground color. This is the color of the lettering.
Background Color	From the drop-down dialog, select a background color. This color fills the footer box.
Text: Title, Subtitle, [other]	
Text	Enter the text you want to use as the title of your report, such Top 10 Attacks per Zone.
Horizontal Alignment	From the drop-down menu, select where you want the title to appear in the title area: left, right or center.
Vertical Alignment	From the drop-down menu, select where you want the title to appear in the title area: top, center, or bottom.
Font	From the drop-down dialog, select a font from the list of fonts available on your local system, font size, and style (bold, italic). The preview window indicates how the font will look.
Foreground Color	From the drop-down dialog, select a foreground color. This is the color of the lettering.
Background Color	From the drop-down dialog, select a background color. This color fills the title box.

Preview

The Preview section shows the layout of the report, and does not show the formatting updates as you go along. If you have designed other text boxes for your template, the attributes for those text boxes are displayed here using the same format as those shown above.

Binding Data to the Report

After the template is chosen and formatted, you are ready to populate the elements of the report with data.



Related topics:

- ["Binding Data to Charts" on the next page](#)
- ["Selecting Data for the X-Axis on a Chart" on page 401](#)
- ["Selecting Data for the Y-Axis on a Chart" on page 403](#)
- ["Selecting Data for the Z-Axis on a Chart \(Optional\)" on page 405](#)
- ["Specifying Top/Bottom Filters Aggregation Filters for a Chart \(Optional\)" on page 407](#)
- ["Setting Display Options and Scale Formats for Charts" on page 409](#)

Binding Data to Charts

The **Data** tab is where you choose which parts of the query or filter result data you want to use for each report element, apply legends and, optionally, top/bottom functions.

Use these options to select the data source (query or trend), chart and table type to use for the report, columns to include, and details on how the chart presents data.

Chart Data	Description
Data Source	From the drop-down menu, select an existing data source you want to use for the chart in your report. The data source drop-down menu provides a list of existing resources based on the resource type you selected in the accompanying drop-down. You can report on queries, trends, active lists, or session lists. When the data source is selected, the remaining elements of the Data tab populate with the data available in the selected resource.
Chart Type	From the drop-down menu, select the type of chart you want to use for the chart part of the report. Depending on the template you use, you may have are several types of bar charts available as well as line charts, pie charts, speedometer, and so forth. The data source and chart type you choose apply to both the X- and Y- axes.
2 sets of data (z-axis)	Check this box if you want to plot the same type of data in a series fashion, for example, hourly EPS rates for connectors and different connectors comprise the series. Note: If you check this box, the Data tab adds a subtab for the Z-axis.

Selecting Data for the X-Axis on a Chart

Data Source: Choose an existing resource on which to run the report. (Available sources will depend on which resource type is selected.) Data selected here apply to both X and Y axes.

Resource Type: Choose the resource type. You can report on queries, trends, active lists, or session lists.

Chart Type: Choose a chart type: Bar, line, pie, and so on.

X, Y Axis Tabs: Indicates which aspect of Data details you are working on. Optionally, you can aggregate on one of the columns for this report.

X-Axis Details: Choose and order the data you want to use on the X-axis. Optionally, create and position a label for the data.

If the report template you selected contains a **Chart**, bind your result data to the chart as described in the following table.

Binding Results to a Chart

X-Axis Data Attribute	Description
Available Columns	Select the data fields from the query you want to show in the X-axis and use the right-hand arrow to move it to the Selected Columns area. The data you select here should be the items you want to count For example, to build a trend report showing number of events over time, use a trend that captures the number of events per day. Add the end time to the X-axis to represent the day and add the count gathered for that day to the Y-axis. In this case, the X-axis is the data label, and the Y-axis is the count.
Selected Columns	The Selected Columns area shows which data fields you have selected for the X-axis, and provides the opportunity to change the sort order of the data.

Binding Results to a Chart, continued

X-Axis Data Attribute	Description
Sort	Sort the selected column in ascending or descending order. To change the sort order, select the column then click the Sort check box. Select A-Z to sort data in ascending order; select Z-A to sort data in descending order. Note: If you are also plotting a Z-axis (data series) and your output is a bar chart, make sure to set your sorting on the X-axis.
X-Axis Title	Specify a title for the X-axis.
Label Rotation	Select a rotation angle for the by entering a digit between 0 and 90. Labels refer to the individual X-axis data points, which are automatically derived from the data. The Label Rotation controls the angle of these labels.

Selecting Data for the X-Axis on a Chart

Data Source: Choose an existing resource on which to run the report. (Available sources will depend on which resource type is selected.) Data selected here apply to both X and Y axes.

Resource Type: Choose the resource type. You can report on queries, trends, active lists, or session lists.

Chart Type: Choose a chart type: Bar, line, pie, and so on.

X, Y Axis Tabs: Indicates which aspect of Data details you are working on. Optionally, you can aggregate on one of the columns for this report.


X-Axis Details: Choose and order the data you want to use on the X axis. Optionally, create and position a label for the data.

If the report template you selected contains a **Chart**, bind your result data to the chart as described in the following table.

Binding Results to a Chart

X-Axis Data Attribute	Description
Available Columns	Select the data fields from the query you want to show in the X-axis and use the right-hand arrow to move it to the Selected Columns area. The data you select here should be the items you want to count For example, to build a trend report showing number of events over time, use a trend that captures the number of events per day. Add the end time to the X-axis to represent the day and add the count gathered for that day to the Y-axis. In this case, the X-axis is the data label, and the Y-axis is the count.
Selected Columns	The Selected Columns area shows which data fields you have selected for the X-axis, and provides the opportunity to change the sort order of the data.
Sort	Sort the selected column in ascending or descending order. To change the sort order, select the column then click the Sort check box. Select A-Z to sort data in ascending order; select Z-A to sort data in descending order. Note: If you are also plotting a Z-axis (data series) and your output is a bar chart, make sure to set your sorting on the X-axis.
X-Axis Title	Specify a title for the X-axis.
Label Rotation	Select a rotation angle for the by entering a digit between 0 and 90. Labels refer to the individual X-axis data points, which are automatically derived from the data. The Label Rotation controls the angle of these labels.

Selecting Data for the Y-Axis on a Chart



Y-axis data should be numeric. If the data you select is a non-numeric data type, such as a string, apply a numeric *summary function* to it, such as **Count** or **Count distinct**.

Y-Axis Attributes for Charts

Y-Axis Data Attribute	Description
Available Columns	Select the data fields from the query you want to show in the Y-axis and use the right-hand arrow to move it to the Selected Columns area. The data you select here should be the item you want to count by. For example, to show how many addresses each of your attacker zones have, you would select the attacker address.
Selected Columns	The Selected Columns area shows which data fields you have selected for the Y-axis, and provides the opportunity to change the sort order of the data.

Y-Axis Attributes for Charts, continued

Y-Axis Data Attribute	Description
Summary Function	You can assign a summary function to one or more columns of data. (In the “Function” row for a column, click in the column to get a drop-down menu of functions.) • Count - Provides a count of all line-items returned by the query. Note: The Count function is a simple count of all events. It takes into consideration the <i>aggregated event count</i> and counts each event in an <i>aggregated event</i> individually. For example, if an event has an aggregated event count of 5, the Count function counts this event as equivalent to 5 events (with an <i>aggregated event count</i> of 1 each). Take this into account when comparing the number of rows in a report with the “grand total” count based on the Count function. • Count Distinct - Provides a count of how many items are unique. For example, if there are 100 IP addresses but only 5 of them are unique, the system counts 5. • Average - Adds the results of numeric data and divides by the number of line items. • Sum - Adds the results of numeric data. • Max - For numeric data, Max calculates the line item with the highest value. • Min - For numeric data, Min calculates the line item with the lowest value. • Median - For numeric data, Median calculates the line item with the value closest to the middle between high and low. • Standard Deviation - For numeric data, measures the dispersion of the values in the data set (how spread out they are). If the data points are all close to the mean, then the standard deviation is close to zero. If many of the data points are far from the mean, then the standard deviation is further from zero. If all the data values are equal, then the standard deviation is zero. The Standard Deviation is the square root of the variance. • Variance - For numeric data, measures how spread out the distribution of data is. The variance is computed as the average squared deviation of each number from its mean. The variance and the standard deviation are closely related measures of dispersion and variability. Selecting one of these functions activates the Aggregation tab, where you can set further parameters on these functions. To set a function, select a column, and choose a function from the Summary Function drop-down menu.
Y-Axis Title	Type in a title for the Y-axis.
Label Rotation	Select a rotation angle for the by entering a digit between 0 and 90. Labels refer to the individual Y-axis data points, which are automatically derived from the data. The Label Rotation controls the angle of these labels.
Sort by	Optionally, choose a sorting order for the data on the Y axis. You can display data alphabetically (the default), reverse alphabetical, or sort by count. Caution: If you are plotting a Z-axis (data series) and your output is a bar chart, make sure to set your sorting on the X-axis.

Selecting Data for the Z-Axis on a Chart (Optional)

If you want to plot data in a series fashion, check the [2 sets of data \(z-axis\)](#) checkbox in the Data tab. This makes the Z-Axis subtab available, where you can select a column for your Z-Axis.




Tip: Best practices when using series (Z-axis):

If you are setting data on the Z-axis, note that you cannot use Top N Charts. Also consider the effect of the [Row Limit](#) setting in the query. For example, if your row limit for a normal chart (X- and Y-axes only) is 10 rows, you will get 10 data points.


For a stacked chart with Z-axis, you will still get 10 data points but two will be in the same column. You can try a limit of 25 to 30 rows, which makes the chart readable, but could cause a readability issue if there is not a lot of overlap (stacking) in the data series.

Effect of Sorting on Bar Charts with Series Data

You might encounter unexpected results in your bar charts (stacked or not stacked) if you are plotting data on the X-, Y-, and Z-axis. Consider the following stacked bar charts, where the expected result is shown in Example A but the actual result shows as Example B. In Example B, the stacks seem to spread out along the X-axis as separate data points (the samples are only to represent the behavior described in this topic).




Example A: Expected Result



Example B: Actual Result


The following configurations can result in a stacked chart report similar to Example B:

- The query has an ORDER BY statement:




The query is not wrong; in fact, the report rendered as a table shows correctly. Additionally, the query may be driving other resources like query viewers and so on; therefore, you should not change the query in this case.

- The chart is configured for z-series data plotting:



Again, there is nothing wrong with this setting if you want to plot a data series.


- The column included in the ORDER BY statement from the query is applied to the Y-axis and sorting is also enabled:



Using the query's ORDER BY column on the Y-axis and also applying the sorting to the Y-axis gives you the Example B report, the one you did not expect. To fix this problem, configure the report so that sorting is ignored on the Y-axis. Instead, apply sorting to the X-axis according to the following instructions.

To configure the report so that it sorts on the X axis:


1. Go to the chart's setting for the X-axis parameters and click the entry under Selected Columns.
2. Click the **Sort** checkbox and select the sort order as desired.



You do not need to remove sorting on the Y-axis as long as sorting is enabled for the X-axis. The X-axis takes precedence, and the resulting report will now be like Example A.

Specifying Top/Bottom Filters Aggregation Filters for a Chart (Optional)

You can also set **Top/Bottom Counts** for a chart. This tab only becomes active when a summary function is applied to data in the Y axis. Settings in the **Aggregation tab** set top/bottom counts to data with summary functions applied. This is an optional step.




On the Chart Aggregation tab, set the top or bottom filter for the chart. If there are more charts in your report, repeat these processes until data is bound to all the charts and laid out in your report template.

Aggregation Top/Bottom Filter	Description
None (Show all)	By default, no top/bottom filter is set.
Top	Select Top if you want to show the a certain number of entries with the highest values. Enter a digit in the text box, and from the drop-down list, select an appropriate Y-axis data column with a function applied.
Bottom	Select Bottom if you want to show a certain number of entries with the lowest values. Enter a digit in the text box, and from the drop-down list, select an appropriate Y-axis data column with a function applied.

Setting Display Options and Scale Formats for Charts

Use the Display Options subtab to specify visual formats on the chart. These options, unless otherwise noted, simultaneously apply to the X- and Y-axes.



Category	Attribute	Description
Display Options		
	Font	From the drop-down menu, select a font style and size for the X- and Y-axis text.
	Show Legend, Placement	Select this box to show a legend of the data elements. Keep in mind the number of different data elements your query may return. If the data you selected contains many elements, the legend may be large, which reduces the available space for the chart itself. If you choose to display the legend, you can move its location from choices in the Placement drop-down menu.
	Show X-Axis grid, Show Y-Axis grid	For the selected axis, include the grids on the report. The X-axis grid displays as vertical lines per data point, and the Y-axis grid displays as horizontal lines.
Scale and Format		
Data (x-axis)	Show Labels	Select this box to display the labels for data points on the X-axis.

Category	Attribute	Description
Data (y-axis)	Logarithmic Scale	Use this option if your data point deltas are so huge that incorporating them on the chart poses a readability issue; for example, one data point is 1,000 and the next is 100,000. With logarithmic scaling, values are then displayed in orders of magnitude instead of linearly.
	Show Value	Displays the value of the data point.
	Min	The minimum value as a starting data point to plot on the chart. For example, if you set it to 1000, values below that will not be plotted. The chart begins at 1000 on the Y-axis. If left blank, plotting starts at the 0,0 position.
	Max	The maximum value for the data point.
	Incr	Increments to use on the Y-axis starting with the specified minimum.

Binding Data to Tables

If the template you selected contains a table, use the Table **Fields** tab to build a visual representation of a table in which to display the query result. You can choose the type of data source (trend, query, active list, or session list) and the particular data source (which query, trend, and so forth) to report on. Then you can select which fields from the data result to appear in your report (with the “Use” check box). Use Groups to combine fields into a single column in your Report table (drag and drop or menu commands available for the report table).

Table Data	Description
Data Source	From the drop-down menu, select an existing data source you want to use for the table part of your report. The data source drop-down menu provides a list of existing resources based on the resource type you selected in the accompanying drop-down. You can report on queries, trends, active lists, or session lists. When the data source is selected, the remaining elements of the Data tab are populated with the data from the selected resource.

Data Source: Choose an existing resource on which to run the report. (Available sources will depend on which resource type is selected.)

Resource Type: Choose the resource type. You can report on queries, trends, active lists, or session lists.

Fields and Groups: Select the data fields you want to display in the table (by enabling **Use** for a field you want to show).

Provide a display name alias and specify data sort order in the column by dragging and dropping fields to desired location in column order.

Group multiple fields in a single column by dragging fields into Groups.

Display Options: Select a column or group to set font, foreground, and background display options.

Global Options: Global options apply to the whole table.

Related topics:

- ["Specifying Fields for a Table" below](#)
- ["Enabling the Aggregation Tab for a Table" on page 414](#)

Specifying Fields for a Table

In the **Available Columns** area, you can select the fields you want to display in the table, group multiple fields into a single column as needed, assign Alias names for column headings, specify a data sort order, and set column size and alignment options.

Available Columns for Tables

Attribute	Description
Groups	Optionally, you can sort data results from queries by grouping two or more fields into a single column. To create a group: Right-click in the Groups row for a column and choose Make Group. In the dialog, enter a name for the new group and add the selected field. To add fields to a group: Drag fields from the Fields row to the Groups row. Alternatively, right-click a field and choose Add to Group. In the dialog, enter a the group's name.
Fields	This displays fields as columns for your report. The field name is displayed as it is referred to in the database. This field is not editable.
Use	By default, all data entries are selected for use in the table. If you do not want to use all the available columns, uncheck the corresponding check box. Caution: If you de-select a data entry because you do not want to use that column in the report, the column is automatically pushed to the far right (the end of the table) to move it out of the way so that you can focus on the columns you are using. If you then select "Use" again for that same data entry, its column is re-inserted into its original position along with the other columns you have selected to use.
Function	To set a function on a field, click in the Function row for that field's column. Select the function you want to apply to the column from the Function drop-down menu. After the function is set, the field appears with the function icon (. Fields set with functions can be filtered in the Aggregation tab.
Alias	Enter a display name alias for the data column. For example, if the column is referred to as Source Translated Zone Name in the ArcSight database, you can shorten this name to Zone Name or Src Zone for display in the report table. In our example, we provide the aliases Time instead of Timestamp and Number of Logins for Category Outcome (Count).
Width	Set column Width to either of the following options: • Auto - Automatically divides column width evenly among the selected columns • User Specified Layout - This option requires that you enter numbers to specify percentage widths for individual columns.
Sort	Indicate the sort order for the data in each column.
H Align	Right-click in the H Align row to display a drop-down menu for specifying horizontal alignment of text in a given column. You can select text to be left- aligned, centered, or right-aligned in the corresponding column.
V Align	Right-click in the V Align row to get a drop-down menu for specifying vertical alignment of text in a given column. You can select for top, bottom, middle, or baseline text in the corresponding column.
Page Break	Right-click in the Page Break row in a column to get options for specifying a page break before or after the that column.

With the **Custom Layout** options, you can specify custom column widths for the data in the table. By default, the **Custom Layout** drop-down menu shows User Specified Layout, which

enables you to enter a numeral to specify a percentage for individual columns. Select one of the following:

- **Fit content** - Adjusts the column width to accommodate its content without wrapping. If the content is wider than the table, the table is extended to multiple pages.
- **Fit content one table area per page** - Adjusts the column width to accommodate its content without wrapping, and breaks each column onto its own page.
- **Fit content to page** - Adjusts the column width to accommodate its content without wrapping, and stretches the last column to fill the page.
- **Equal width columns** - Each column receives the same width to fit across a single page.
- **User specified layout** - Enables you to enter a numeral that represents a percentage of the overall page width. You can set a percentage for each column that totals 100%, or enter a percentage for one column, and the others selected receive an even percentage of the space remaining.

The **Display Options** area provides format options for each individual data column. This enables you to set different font style, size, and color and column background colors for each data column. To activate the display options, select one or more data columns:

- **To select one column:** click the column by its Field name.
- **To select one or more contiguous columns:** click a field, hold down the Shift key, and select the remaining fields.
- **To select one or more non-contiguous columns:** click a field, hold down the Ctrl key and select the remaining fields

Attribute	Description
Font	From the drop-down menu, choose a font for the selected columns.
Foreground Color	Foreground color for text, any visible lines that describe rows/columns, and other elements in the foreground.
Background Color	Background (field) color for the data column.

In the **Global Options** area, you can set formatting options that apply to the whole table (not just one column).

Attribute	Description
Merge cells	Indicates whether to merge cells for grouped columns. When this option is enabled, identical values in grouped columns show only once. When this option is disabled, identical values show as many times as they occur (regardless of whether they are grouped).
Show group header	Indicates whether to show a group header row. This is a group label for when you have a summary function that adds one more rows at the end of the section. If this option is enabled, the table includes an extra column with a header derived from the content by which the section is grouped.
Show group columns	Enable this option to populate the grouped columns with data. (If this option is disabled, grouped columns are empty.)
Grand total	If you want to provide a grand total of all the sections, check the Show grand total box.
Label	If you selected a grand total, you can apply a label for the grand total. For example, use Total VPN Login Attempts.

Click the **Preview** button to preview the report table with the current configuration.

Enabling the Aggregation Tab for a Table

If your report is using the Table template, the Aggregation tab is disabled by default until the following requirements are met on the report's Data > Fields tab:


- At least one of the columns is assigned to a group.
- At least one of the columns must be set with a function.

The following scenario describes the process to enable the Aggregation tab.

1. Define the query. For example, you want to look at event names by priority. In the query's Fields tab, select the columns you need (Name and Priority in this case). After selecting the columns, add these same columns to the GROUP BY list. In the Conditions tab, define the conditions for the query (for example, you are interested in Priority greater than 3).
2. Create a report. On the Template tab, use one of the table templates (for example, Simple Table Portrait). In the Data tab, specify the query defined in the beginning of this procedure as your data source. The Fields tab is automatically populated by the columns from the query, in this case, Name and Priority.
3. On the Fields tab, create a group for one of the columns (for example, for Priority), then select a function for another column (for example, for Name) from the drop-down list. Columns with functions can be filtered through the Aggregation tab. For instructions on how to add columns to groups and how to apply functions to columns, see ["Query Fields" on page 324](#). The Aggregation tab is enabled after you click **Apply**.

Setting Top/Bottom Counts in Table Aggregation Tab (Optional)

If the Aggregation tab is enabled, you can set optional top/bottom counts to data with functions applied to individual fields. The following example shows the format used is *Function (Field)*. This is an optional step.




On the Aggregation tab, set the top or bottom filter for the table. If there are more tables in your report, repeat these processes until data is bound to all the tables and laid out in your report template.

Aggregation Top/Bottom Filter	Description
None (Show all)	By default, no top/bottom filter is set.
Top	Select Top if you want to show the a certain number of entries with the highest values. Enter a digit in the text box, and from the drop-down list, select an appropriate Y-axis data column with a function applied.
Bottom	Select Bottom if you want to show a certain number of entries with the lowest values. Enter a digit in the text box, and from the drop-down list, select an appropriate Y-axis data column with a function applied.

Setting Default and Custom Report Parameters

The report's **Parameters** tab is where you set report output details such as file format, paper size, and routing instructions. From here you can also set limits on the query return such as row limits and time zone restraints, apply filters, and specify report start and end times.



In the **Report Parameters** area, enter the following values. The Use Default checkboxes do not apply to these items; the default values are used until you uncheck the boxes and use a new value. Note that users can re-set most of these parameters at report runtime.

Report Parameter Values

Common Parameters	Description
Report Format	From the drop-down menu, select one of the following report output formats: • pdf - Outputs the report as an Adobe PDF file. • xls - Generates a Microsoft Excel file for tables and charts. Note: XLS reports you run with <i>Microsoft Excel 2002</i> might have page break format problems (misalignments, column spillover) due to default page size settings in Excel. To correct this problem, open the resulting XLS report in Excel, choose File > Page Setup from the menus, change the paper size to Letter (instead of Legal), and click OK to save your changes. The report has the appropriate page break formatting. <i>This problem does not occur in newer versions of Microsoft Excel.</i> Note: XLS report formats display speedometer charts as pie charts. This is a known limitation in Microsoft Excel. • rtf - Produces a rich-text format document. • csv - Creates tabular data as a list of comma-separated values. Note: Reports generated in CSV format are not the full equivalent of exports to other formats like PDF or HTML. CSV format is useful for loading report data into a spreadsheet for further manipulation. Since CSV is meant to contain tabular data, only the table data of a report is normally useful. Therefore, ArcSight exports only the table data portion of a report to CSV format, ignoring any other report information such as charts or text, including report titles. • html - Generates the report in a Web page displayed by the default web browser. Note: Running a report with a huge row limit (10,000 rows) in HTML format may cause the browser to be unresponsive. In such cases, run the report in PDF. Your selection affects your choice for e-mail formats. See description for Email Format, another report parameter, in this table. The output file is created at report runtime and is stored in the Archived Report Group corresponding to the report's group. See the Archives tab on the report's edit panel.
Page Size	From the drop-down menu, select a paper size.
Run as User	This option is not available for case reports. Run the report as a particular user. From the drop-down menu, select the user name by which you would like to run the report. For example, this option would allow an administrator for a Managed Security Service Provider (MSSP) to run report for a customer. The administrator would need write permissions to the user.

Report Parameter Values, continued

Common Parameters	Description
Email to	You can have the report sent as e-mail to one or more ArcSight users. From the drop-down menu, select the Console users to whom the report should be e-mailed. The selection list is read from the Users resource). The recipient will only see his or her user name in the To field even if there are multiple recipients for this report. Handling empty reports: If you are emailing reports, empty reports will also be sent. This is determined by the server property, <code>report.scheduler.notify_empty_reports</code>, which is set to true. If you do not want empty reports to be sent, add the property to the <code>server.properties</code> file and change the setting to false. Follow the instructions in the <i>ESM Administrator's Guide</i> on how to edit this file. The details are in the guide's topic, "Managing and Changing Properties File Settings."
Email addresses	Send the report to one or more comma-separated or semicolon-separated e-mail addresses. This option does not require the recipient to be an ArcSight user. Note: The recipient will only see his or her e-mail address in the To field even if there are multiple recipients for this report.
Email Format	Specify how the report is to be accessed by the recipient. • Select Send URL if you want to point users to the report. Use this option if the report is large and is saved (archived) to a network-accessible location You can provide URLs for all report formats: PDF, XLS, RTF, CSV, and HTML. • Select Attach Report if you want to send the report directly to the user's e-mail box. You can only attach PDF, XLS, RTF, and CSV report formats. • Select Attach Compressed Report if you want the PDF, XLS, RTF, or CSV report to be compressed (zipped) first before mailing. Note: The notification automatically uses the URL if you select an email format for an unsupported report format, or if the attachment is too large to be sent. You can adjust the maximum attachment size up to 100 MB with the property, <code>report.upload.maxFileSize</code>. Set this property to the size you want in the <code>server.properties</code> file. Refer to the <i>ESM Administrator's Guide</i> for procedures to edit properties files.
Email Subject	Specify the subject on the notification. Defaults to the report's Name attribute (denoted by <code>\$ReportName</code>). If you want to use a customized subject, type the text either in addition to the default or replace the default text entirely.



Note: If you are emailing reports, empty reports will also be sent. This is determined by the server property, `report.scheduler.notify_empty_reports`, which is set to **true**. If you don't want empty reports to be sent, add the property to the `server.properties` file and change the setting to **false**. Follow the instructions in the *ESM Administrator's Guide* on how to edit this file. The details are in the guide's *Configuration* section, topic on *Managing and Changing Properties File Settings*.

Adding Custom Parameters for Report Data

Purpose:

Custom parameters override the parameters defined in the query. Custom parameters are saved locally to the report definition, and are not persisted back in the query.

For example, if you want all the report elements to report on events for the past 2 hours, you can create a start time parameter of `$Now-2h`, which sets both table and chart start times to `$Now-2h`.

To add a custom parameter that applies to the Report data:

1. On the report's Parameters tab, click the **Add** button.
2. Give the parameter a name and map it to a query parameter.
3. Click **OK** to apply them to the report definition.
4. Back in the **Parameters** tab in the **Custom Parameters** section, enter an override parameter for the fields you selected from the Add Custom Parameters dialog.
5. In the **Query Parameters** area, enter any override values for the parameters in your query data. The Use Default checkboxes are only activated for items where default parameters exist and override values can be entered.

Enter these override parameters as needed for each chart and table.

Query Override Parameters	Description
Time Zone	By default, the Manager Time Zone is used. Choose the Console time zone, or another time zone from the drop-down list.
Filter By	Set a filter to operate on the query conditions.

Query Override Parameters	Description
Row Limit	The report gets its row limits from the settings in the query being used. By default, row limit for a table is 10000 and row limit for a chart is 25. You can change the default to manage row size.
Start Time	The report gets its start time from the query being used. To override the one set in the query, disable Use Default for this field and specify a start time here. For example, if you want all the report elements to report on events for the past 2 hours, you can create a start-time parameter of <i>\$Now-2h</i> , which sets both table and chart start times to <i>\$Now-2h</i> . This setting is saved locally as part of the report definition, not as part of the original query upon which the report is based.
End Time	The report gets its end time from the query being used. To override the one set in the query, disable Use Default for this field and specify an end time here. This setting is saved locally as part of the report definition, not as part of the original query or trend upon which the report is based.

- Click **Apply** to save settings or **OK** to save settings and close the Inspect/Edit details for this report.

Displaying a Custom Parameter Prompt at Report Runtime

You can configure your report to prompt for a value at report runtime. If so, the user is prompted by the Parameters dialog box to enter a value before the report is created. This prompt will be based on one of the report's custom parameters. For example, you have a report which is based on a query on assets, and you want the user to enter the first few characters of the host name when the report is run, so that the report contains data only about those hosts. The configuration is a two-step process:

- ["Defining the Prompt in the Query's Condition Tab" on page 423](#)
- ["Adding or Removing a Prompt for Custom Parameters in the Report" below](#)


Adding or Removing a Prompt for Custom Parameters in the Report

After following the instructions in ["Setting Default and Custom Report Parameters" on page 415](#), follow the instructions here to add the prompt to the report's parameters. Instructions are also provided for deleting custom prompts if you no longer require them.

To add the prompt:

1. On the Reports resource, select the **Reports** tab.
2. Locate the report, right-click, and choose **Edit Report**.
3. Go to the **Parameters** tab and verify that the field configured in the base query is listed in the Query Parameters list.
4. Click the **Add** button ( Add...) to display the Add Custom Parameter dialog.
5. Under Map to Parameters, check the parameter corresponding to the prompt name, then enter a meaningful name in the Name field.

The name used here will be displayed to the person running the report, so make sure the name you use clearly states what the report is expected to display. For example, you want host names to start with a specified text string, as shown in the Add Custom Parameter dialog:



Caution: In the query condition, spaces are resolved with underscores but not here, so make sure you insert the underscores as required.

Click **OK**.

The custom parameter is added to the Report Parameters section.

Name	Value	Use Default
Common Parameters		
Report Format	pdf	☐
Page Size	Letter [8.5x11 in]	☐
Run as User	Select a User	☐
Email to		☐
Email addresses		☐
Email Format	Send URL	☐
Email Subject	\$ReportName	☐
Custom Parameters		
Hostname_Starts_with		☐

Name	Value	Use Default
Table		
Time Zone	Manager Time Zone	☐
Hostname_Starts_with	1/4	☒
Row Limit	10000	☒

Notice that the default value for the custom parameter in the Query Parameters section is not picked up in the Report Parameters section. If you want the default value to be displayed by the prompt, enter the value in the Report Parameters section.

6. Click **Apply** or **OK** to save the report.

The next time you run the report, the custom parameter will be displayed in the Report Parameters dialog as in the following example:

Name	Value
Common Parameters	
Report Format	pdf
Page Size	Letter [8.5x11 in]
Run as User	Select a User
Email to	
Email addresses	
Email Format	Send URL
Email Subject	\$ReportName
Custom Parameters	
Hostname_Starts_with	
Timezone	Manager Time Zone
RowLimit	

☐ Save Output

OK Cancel Help

Based on the example, you can then enter the starting character strings (case-sensitive) of host names to be included in the report. If a default value is displayed, choose to run the report with the default or replace the value before running. See ["Running and Managing Reports" on page 446](#) for more details on how to run a report.

To remove the prompt:


If the report no longer requires a prompt, revert your prompt configurations in this sequence:

1. On the report's **Parameters** tab, select the custom parameter.
2. Click **Remove**.
3. On the report's **Queries** tab, go to the **Conditions** tab.
4. Right-click the condition for the prompt and select **Delete**.

Defining the Prompt in the Query's Condition Tab

The following procedure instructs you to create a query condition to associate a prompt with a specific field used by the query.

1. On the Reports resource, select the **Queries** tab.
2. Locate the query being used by the report, right-click, and choose **Edit Query**. This example uses an asset-based query.
3. Go to the query's **Conditions** tab.
4. On the Field Set or CCE panel at the bottom of the tab, locate the field for which you want to display a prompt, for example, **Host Name**.
5. To prompt for a value for the Host Name field:
 - a. Select a logical operator. This example uses **StartsWith**.
 - b. Click the **Browse** icon ().
 - c. In the Advanced Editor dialog, click the **Parameter** check box to display the Prompt and Default Value fields.
 - d. In the Prompt field, provide a meaningful name for this prompt. For example, your prompt can say **HostName Starts with** to match your logical operator. Notice that spaces are automatically replaced by underscores.
 - e. Enter a default value (case sensitive) which can be changed at report runtime. The following example on the Advanced Editor dialog uses **VA** for host names starting with that string.






Note: When specifying default values for the prompt, remember that string values are case sensitive.

- f. Click **OK**. Click on the field again on the Field Set or CCE panel to save your condition.

The Console automatically inserts the **@** symbol in front of your prompt name, replaces spaces with underscores, and also adds this condition statement as shown:



The @ symbol is automatically inserted and spaces are replaced by underscores.

6. Save your query. The condition you just added will be displayed on the report's query parameters list. The next step is to add this prompt as a custom parameter in the report.

Generating Reports with Central European, Cyrillic, or Asian Fonts

This topic provides setup procedures so reports can display Central European, Cyrillic, or Asian (Chinese, Japanese, or Korean) character sets.

To generate reports that properly display Central European, Cyrillic, or Asian character sets:

1. Configure the operating system and the ArcSight Manager to support the Language you are using.
2. Make sure you have the Adobe Acrobat Reader 9 or later to view the PDF report.

If the Manager is running on Linux, do the following:

1. Download ARIALUNI.TTF font from the Linux support site.
2. Go to the /usr/share/fonts/ directory and create a subdirectory called /arial.
3. Copy ARIALUNI.TTF to /usr/share/fonts/arial.
4. Make a backup of the \$ARCSIGHT_HOME/reports/sree.properties file.
5. Add this property to sree.properties:
font.trueType.path=/usr/share/fonts/arial.

To generate a report in PDF format to display Central European, Cyrillic, or Asian characters:

1. Log in to the ArcSight Console and open the report.
2. Find the template used by the report.
3. Edit the template and select Open in Designer.
4. Edit the fields that need to display these characters.
5. Set the fonts to **Arial Unicode** for the fields that display these characters
6. Save the template and click **Apply**.
7. Run the report with PDF format.
8. Open the generated report (using Adobe Acrobat Reader 9 or later for PDF) to see these characters.


To generate a report in RTF format to display Central European, Cyrillic, or Asian characters:

1. Log in to the ArcSight Console.
2. Select **Edit > Preferences > Global Options**.
3. Set the font to **Arial Unicode MS**.

Creating Focused Reports

In addition to using the reports already available in the Navigator panel's Reports resource tree, you can easily make and save refinements to these definitions. These more narrowly defined or focused reports are also stored in the resource tree, so other people can also use them.


Focused reports are identical to other reports. They differ only in being useful variations on already defined reports. You create focused reports when you want to make a special variation available to other ArcSight users through the Reports resource tree.



Where: Navigator > Resources > Reports

To create a focused report:

1. Right-click the report and select **New Focused Report**.
2. In the Focused Report Editor, select the **Attributes** tab and name the report. Name focused reports in a fashion that properly distinguishes them from their originals.
3. Click the **Parameters** tab and change any of the values as appropriate. These values are the same ones you set when Running a New or Archived Report.



4. Click **Apply** to make changes and keep the editor open. Click **OK** to store the definition in the resource tree in the same folder as the original report and close the editor.

Using Report Templates

Template definitions determine how query and trend data are displayed in a report. You can edit ArcSight-provided report templates or create your own templates using the Template Designer Wizard.

A template consists of report design elements, such as headers, footers, title bars, charts, and tables, arranged on a page according to a layout specification. Templates can accommodate input from multiple queries and show multiple visual elements, such as three charts and a table each pulling from a different data source, in a single report.

Use the designer wizard to create and adjust templates to specify which data is displayed, what visual elements are used (variations on tables, charts, graphs, and so on), the layout of those elements, the report output file format, and much more.

To navigate to the report templates:

In the Navigator panel, select **Reports** resource from the drop-down menu and click the **Templates** tab.

Report templates are a component of ArcSight Reporting resource tools.

Related topics:


- ["Applying a Standard Template to an Existing Report" below](#)
- ["Creating a New Report Based on a Template" on the next page](#)
- ["Copying a Standard Template" on page 429](#)
- ["Editing a Template" on page 430](#)
- ["Building Reports" on page 388](#)

Applying a Standard Template to an Existing Report

Where: Navigator > Reports > Reports tab

1. Double-click the report to which you want to apply a template. Alternatively, you can select the report, right-click and select **Edit Report** from the context menu.
This brings up the Report editor in the Inspect/Edit panel.
2. In the Report editor, click the **Template** tab for the selected report.
3. In the **Report Template** field drop-down menu, select a template.
4. Click **OK** to apply the template and close the file browser.

5. Click **OK**.



Creating a New Report Based on a Template

Where: Navigator > Reports > Templates tab

1. Right-click the template you want to use and select **New Report from Template**. This launches the Reports Editor in the Inspect/Edit panel with the chosen template.
2. See ["Creating or Editing a Report" on page 393](#) for details on how to define data for your report and fine-tune the template by means of the Template tab in the Report editor for this report.

Copying a Standard Template

An easy way to start customizing a template is to copy an existing template and modify it to suit your needs.

Where: Navigator > Resources > Reports > Templates tab

To copy a template:

1. Select the standard template you want to copy.
2. Left-click, and drag and drop the selected template into your user folder.
3. Select **Copy** from the Drag & Drop Options dialog. A copy of the template is dropped into your user folder.

Alternatively, you can select the template you want to copy in the Navigator and choose **Edit > Copy** from the menus. Then select your user folder and click paste to drop the template into the folder.

Managing Report Template Groups

This procedure instructs you to create report template groups (folders) to store custom templates. Report template groups are like any other ESM resource group. By default, if you create your own templates, the Console adds them in your user folder in the Reports > Templates tab.

Where: Navigator > Resources > Reports > Templates tab

To create a report template group:

1. Select the parent template group for your custom template group.
2. Go to ["Adding or Editing a Resource Group" on page 711](#) for related steps.



Caution: For the group's name, do not use special characters like & (ampersand), * (asterisk), and % (percent). The group name becomes part of URIs (uniform resource identifiers), and therefore follows the same restrictions for URIs in general. The escape character option is not available in this case.

3. Click **Apply** or **OK**.

Editing a Template

The Report Designer is powered by InetSoft, which also provides the Report Designer's online help. Refer to the Report Designer's online help for procedures on how to customize your own template.

1. Select the **Reports** resource in the Navigator.
2. Click the **Template** tab.
3. Right-click a template and choose **Launch Designer**, or choose **Edit Template** and click the **Open in Designer** button on the Attributes tab for the template editor.
4. Refer to the documentation by selecting **Help** from the designer toolbar.

End-to-End Reporting Examples

This topic includes two examples:

Quick-start example with Report Wizard - An introductory example of how to create a simple report on the results of a single, stock query with the Report Wizard.

Advanced example - A more in-depth example reporting on the results of several trend-queries and using a heavily-modified 3-charts template. This example walks you through creating the following resources for example queries, trend, and report:


- A base query that captures data about number of VPN login attempts per hour.
- A trend that takes the base query as input, executes it, and stores captured data per a schedule you define.
- Queries that build on the trend to filter on various VPN login outcomes.
- A report that uses the complex queries as data sources and provides visual representations of query results in charts and tables based on an ArcSight provided template.

Even if you do not anticipate immediately having to create these elements from scratch (ArcSight provides a starter set of stock reporting content), we suggest working through both the simple example and the more complex one to gain an understanding of how queries, trends, and templates work together in the context of reporting.

Example of Creating a Simple Report with the Wizard


Where: Navigator > Resources > Reports

1. Right-click your Reports user folder and select **Start Report Wizard**.
Click **Next** on the wizard's welcome page.
2. On the **Report Name** page, enter a name for your report (this example uses My Top 10 Events Report). Click **Next**.
3. On the **Data Sources** page, select the **Queries** tab (if not already selected, and navigate the Queries tree to choose an existing query. For this example, we select the **Top 10 Events query**, which you can find in Queries/Shared/All Queries/ArcSight Administration/ESM/System Health/Events/Top N Activity Reports/.



Click **Next**.

4. On the **Template** page, select a template. For this example, select the **Simple Table Portrait** template under /Report Templates Shared/All Report Templates/ArcSight System/.




5. On the **Bind** page, verify the data source you selected in a previous step. In our example, we chose the Top 10 Events query.


Click **Next**.

6. Review the report configuration summary and verify that all previous settings are reflected here.
7. When you are satisfied with the report configuration, click **Finish** on the last page of the Report Wizard.

The **Report Editor** is automatically displayed on the Attributes tab.



8. Click **Apply** or **OK** on the Report editor to apply the report name and create the report.
9. The new report is added to your Reports folder shown in the Navigator.



10. On the Navigator panel Reports tree, open your Reports folder, right-click the new report and select **Run > Report with defaults**.

Advanced Reporting Example Overview

We build an example query that shows the number of login attempts on a virtual private network (VPN). Then, we use the query in a trend to collect data on VPN login attempts on an hourly basis. Next, we build several more focused queries on top of the trend to get views into particular slices of the data (all login attempts, successful logins, and failed logins).

Finally, we use the data results from the queries and trends to create a report. To format the report, we use one of the ArcSight provided templates.

Start by navigating to the Reports resource in the Navigator panel, then follow these steps to build the example report:



Note: You need a set of canned VPN login events to properly verify the query and trend resources created for this example.


Step 1 - Build the VPN Logins Outcome Query

Start by building a base query that captures VPN Login Data to return a count of hourly VPN login attempts. Following is a summary of configuration details you can use to create this query. (If you need more general help on creating queries, refer to Building Queries.)

Query Name and Other General Attributes


Create a new query, name it, and set the required attributes for it on the Query **General** tab as shown.

Query Attributes	Value
Name	VPN Logins Outcome - Hourly
Query on	Event
Start Time	\$Now - 1d
End Time	\$Now
Use as Timestamp	End Time
Row Limit	10000 (this is the default)




Fields to Include in Query Result

On the Query **Fields** tab, select fields and apply functions as shown to populate columns in the table of result data.

Selected Fields	Description
Category Outcome	Add the Category Outcome field to the Query Structure list by clicking the Add SELECT columns link. This enables the Query Structure's toolbar icons and also opens the SELECT panel.  From the Field Sets panel at the bottom, select Category\Category Outcome to add it to the list. This column contains the outcome of each login attempt (success or failure).
Category Outcome (COUNT)	Add a second Category Column to the Query Structure list to which you will apply the Count function. To get this, click the Category Outcome column you just added, then click the Duplicate Column icon on the toolbar. While this duplicate is still selected, click on it to display the column's edit panel. Select COUNT from the Function drop-down menu (the first field), then click the checkmark icon to apply your changes.  Move this column to the top of the list using the arrow on the toolbar.
Hour	Add the variable called Hour which returns the hour value based on the end time of the event. To get this variable, on the SELECT panel click the Fields & Global Variables tab. Expand /All Fields/ArcSight Foundation/Variables Library/Timestamp Formats and select Hour. This column contains the date and time of the login attempt.

After you click **Apply**, the Group By list is automatically populated with Category Outcome and Hour:



Query Conditions

On the Query **Conditions** tab, define some logical conditions for the login data that narrow the query result to return only the data you are interested in.

- Filter on VPN Logins by specifying that each login attempt must be categorized in a specific event category and device group:


Category Behavior = /Authentication/Verify

Category Device Group = /VPN

- Each login attempt must have a target user name value:

Target User Name Is NOT NULL

Here is how the Field Conditions on this query should look in the display:




Click **Apply** or **OK** in the Query Editor to save the new query.


Step 2 - Build the VPN Logins Outcome Hourly Trend

Create a new trend, name it, and set general attributes for it on the trend **Attributes** tab as shown. This trend uses the data results from the VPN Logins Outcome Query you just created. Keep the defaults for trend interval (1 hour to collect data on an hourly basis) and row limit at 1,000 (it stops collecting data when the table is filled at that limit).

Trend Attributes	Value
Name	VPN Logins Outcome - Hourly
Query	VPN Logins Outcome - Hourly
Enabled	On
Trend Interval	1 day
Row Limit	1000 (default)
Offset Collection By	1 day
Trend Schedule	
Schedule Frequency	Daily
Occurs at	Time when trend is being created in HH:MM:SS



Under Data Fields, you can see the fields the trend is getting from the query initially reflected with the original field names: TimeStamp, Count(Category Outcome), Category Outcome, and Hour. For readability, change the first two fields to the aliases Time and Number of Logins as shown below (double-click on the actual field name and type over the existing name).



From here, you can test the trend to ensure you are getting correct data. To do this, click the **Test** button on the trend **Attributes** tab. The **Test Trend** dialog returns an example result set.


For each row, the Trend should return Timestamp, count of login attempts, Category Outcome (Attempt or Failure), Hour (from the Hour variable).

Trends also have schedules. On the Trend **Schedule** tab, define a schedule that specifies how often you want to run the trend. For the example, define this one to run every hour on the hour (Hourly, every 1 hour at “0 minutes after”).

A trend's range defines when to start and terminate the data collection.

The trend starts as you specified and keep going until it is manually terminated.

Here is the data collected from a trend that ran hourly for a few days. You can view result data from your trend in the grid view by selecting the trend in the Navigator and clicking the Data Viewer for it in the right-click menu.



The screenshot shows a window titled "Viewer" with a tab labeled "VPN Logins Outcome - Hourly Details". Below the tab, it says "Name: VPN Logins Outcome - Hourly" and "Filter: No Filter". A status bar indicates "871 shown / 871 matches". The main area contains a table with the following data:

Time	Number of Logins	Category Outcome	Hour
27 Jul 2016 12:48:26 PDT	26	/Attempt	12
27 Jul 2016 12:48:26 PDT	100	/Attempt	13
27 Jul 2016 12:48:26 PDT	12	/Failure	12
27 Jul 2016 12:48:26 PDT	47	/Failure	13
27 Jul 2016 12:48:26 PDT	13	/Success	12
27 Jul 2016 12:48:26 PDT	54	/Success	13
28 Jul 2016 08:48:26 PDT	22	/Attempt	8
28 Jul 2016 08:48:26 PDT	80	/Attempt	9
28 Jul 2016 08:48:26 PDT	10	/Failure	8
28 Jul 2016 08:48:26 PDT	38	/Failure	9
28 Jul 2016 08:48:26 PDT	12	/Success	8
28 Jul 2016 08:48:26 PDT	42	/Success	9
28 Jul 2016 09:48:26 PDT	23	/Attempt	9
28 Jul 2016 09:48:26 PDT	78	/Attempt	10
28 Jul 2016 09:48:26 PDT	13	/Failure	9
28 Jul 2016 09:48:26 PDT	35	/Failure	10
28 Jul 2016 09:48:26 PDT	11	/Success	9
28 Jul 2016 09:48:26 PDT	42	/Success	10

When you are satisfied that the trend is set up correctly, click **Apply** or **OK** in the trend Editor to save the trend.

Step 3 - Filter the Trend Data (Login Attempts, Successes, Failures)


You can further refine the VPN login query data by creating separate queries based on the trend, each of which capture information a particular aspect of VPN login events. Developing several trend-based queries like this (to show different data slices of common scenarios), gives you a rich set of data views from which to run reports later.

Create three more queries all of which use the original trend as their data source, and then further filter the data to show only attempts, failures, or successes, respectively. Use each of these queries, **Attempt**, **Failure**, and **Success**, to further filter the login data captured in the trend:


- **Login Outcome Trend Query - Attempt**
- **Login Outcome Trend Query - Failure**
- **Login Outcome Trend Query - Success**

As an example of how this is done, here are the details for creating one of these; the **Failure** Query definition.

1. Create a query and name it **Login Outcome Trend Query - Failure**.
2. As the query's data source type (Query On), select **Trend** and select the **VPN Logins Outcome - Hourly** trend.
3. In the Query **Fields** tab, select the same fields as in the original query to populate columns.



4. On the **Conditions** tab, specify **Category Outcome = /Failure**. The query only returns the login attempts that failed.



5. Save your changes. You have now built a query that reports on failed VPN login trends.
6. Create the other two queries:
Login Outcome Trend Query - Attempt

Login Outcome Trend Query - Success

Follow the same procedure, specifying the appropriate Category Outcome condition for each.

You are ready to report on the trend data.

Step 4 - Create the VPN Logins Outcome Report on Trend Data


You can leverage multiple data sources in your report. For this example, you can use all three of the VPN Login trend-based queries you just built to create a report.

On the **Reports** tab, create a new report in your user folder and name it **VPN Login Outcome Trend**.

Choose a Template and Bind it to Result Data

A Template defines the visual constructs of a report such as layout, portrait or landscape, number of tables, number and types of charts, placeholders for text areas, and so on. You can find the ArcSight provided templates under **Report Templates/Shared/All Report Templates/Arcsight System/**.

In the Editor (Inspect/Edit panel) for your new report, click the **Template** tab and select the **Three Charts Landscape** to use as the Report Template. (Look in the drop-down tree under 3 charts/Without Table/ to find this template). In the preview panel you can see what the report template looks like. Double-click the template preview to open it in the viewer. Here you can see what the report looks like before adding the data.



On the Reports **Data** tab, you can bind each of the three charts in the template to each of the VPN login “trend” queries. (The data source type for each of these charts is a query, but remember that each of the queries uses a trend as its data source, which, in turn, was built on our original query.)

Chart	Description
Chart 1	- On the Report Data Chart 1 tab, select Login Outcome Trend Query - Attempt as the data source for the first chart. This query returns the number of login attempts over the last hour. - For Chart Type, select a line chart. - On the Data X-Axis tab, add the Hour value to the Selected Columns. This displays the Hour value on the horizontal axis. - On the Data Y-Axis tab, place the Number of Logins (Category Outcome with “Count” applied to it) in Selected Columns. This displays the count of login attempts on the vertical axis.
Chart 2	On the Report Data Chart 2 tab, select Login Outcome Trend Query - Failure as the data source for the second chart. This query returns the number of failed logins per hour. Configure this chart also to show the Hour value on the X (horizontal) axis, and the number of failed logins on the Y (vertical) axis.

Chart	Description
Chart 3	On the Report Data Chart 3 tab, select Login Outcome Trend Query - Success as the data source for the third chart. This Query returns the number of successful logins per hour. Specify the same assignments as the other charts for X and Y axis.



Tip: You can click **Apply** to save the new Report and then continue working. It is a good idea to save frequently.

Use Custom Parameters

On the Report **Parameters** tab, you can view all the common parameters for the report (in Report parameters area), and all the parameters required for each chart (in Query Parameters area).


You can also provide Custom parameters. You can use Custom parameters to tie together similar parameters from multiple queries for one consistent value. For example, we could do this with Start Time and End Time.

Use Custom Parameters to tie together similar parameters (like Start Time and End Time) from multiple queries for one consistent value.

Query Parameters		
Name	Value	Use Default
Chart3		
Time Zone	Manager Time Zone	☐
Row Limit	25	☒
Start Time	\$Now - 1d	☒
End Time	\$Now	☒
Chart2		
Time Zone	Manager Time Zone	☐
Row Limit	25	☒
Start Time	\$Now - 1d	☒
End Time	\$Now	☒
Chart1		
Time Zone	Manager Time Zone	☐
Row Limit	25	☒
Start Time	\$Now - 1d	☒
End Time	\$Now	☒

Create a new Custom parameter called “start_time”:

Click the **Add** button on the **Parameters** tab, and create a new parameter called start_time to prompt for Start Time field values. Map it to “Start Time” for all three charts (Chart 1, Chart 2, and Chart 3).



The custom parameter is added to the list of report parameters under Custom Parameters.

Similarly, add an End Time by adding a new parameter called “end_time” and map it to End Time for all three charts.

On the Parameters tab under Custom Parameters, use the drop down menus to choose the following values for your new parameters:

- Set **start_time** to \$Now-1d
- Set **end_time** to \$Now

Report : VPN Login Outcome Tre

Attributes | Template | Data | Parameters | Jobs | Notes

Report Parameters

Name	Value	Use Default
Common Parameters		
Report Format	pdf	☐
Page Size	Letter [8.5x11 in]	☐
Run as User	Select a User	☐
Email to		☐
Email addresses		☐
Email Format	Send URL	☐
Email Subject	\$ReportName	☐
Custom Parameters		
StartTime	\$Now - 1d	☐

+ Add... Edit... Remove

Query Parameters

Name	Value	Use Default
Chart3		
Time Zone	Manager Time Zone	☐
Filter by	Select a Filter	☐
Row Limit	25	☒
Start Time	\$Now - 1d	☒
End Time	\$Now	☒
Chart2		
Time Zone	Manager Time Zone	☐
Filter by	Select a Filter	☐
Row Limit	25	☒

Preview... OK Cancel Apply Help

Click **Apply** or **OK** in the Report Editor to save the new report.

Step 5 - Run the Report


Where: Navigator > Resources > Reports

To run the report:

1. Right-click the VPN Login Outcome Trend report.
2. Select **Run > Report with defaults**.

The report will now have three charts each showing a different slice of the data:

- Number of login attempts
- Number of failed logins
- Number of successful logins



Chapter 14: Running and Managing Reports

This topic describes how run and manage various types of reports. Information is included on working with ad hoc reports, archived reports, focused reports, delta reports, and scheduled reports. The topic on Running Reports includes detail on setting report parameters at run-time. Also included is information on how to import and export reports, and work with report groups.

Defined reports are usually run on a schedule and their output archived automatically. But there are also many occasions when you need to run the basic report types directly.

See also ["Building Reports" on page 388](#) for an overview of all reporting tasks and tools, including how to develop new reports, queries, or trends using a provided template.



Tip: Best practices

- By default, you can run up to four reports at the same time. The number of reports allowed to run simultaneously is a configurable parameter on the Manager in ARCSIGHT_HOME/config/server.properties through the setting

```
report.maxconcurrentinteractivereports = <number>
```

If the number is exceeded, each report will be placed in queue and will be run in the order they appear. See "Managing and Changing Properties File Settings" in the [ESM Administrator's Guide](#).

- If you are having problems running a large or complex report, refer to the topic ["Running Large or Complex Reports" on page 456](#).
- If you are having problems running PDF reports with Asian fonts, see the topic ["Generating Reports with Central European, Cyrillic, or Asian Fonts " on page 424](#).
- If a report does not show up as expected, try restarting the Console and running the report again.

Running a Report

When you run reports, you most often use an existing report definition; or a copy of a report already defined, run, and archived for later use. Defining new reports is a separate topic described in Creating a Report. Please see also ["Archiving a Report" on page 459](#) and ["Scheduling Report Tasks" on page 461](#).



Tip: If you are having problems running a large or complex report, refer to the topic ["Running Large or Complex Reports" on page 456](#).

If you are having problems running PDF reports with Asian fonts, see the topic ["Generating Reports with Central European, Cyrillic, or Asian Fonts " on page 424](#).

Where: Navigator > Resources > Reports

To run a defined report:

1. Select the report you want to run.
2. Right-click and select **Run** with one of the report-type options described below:

Report Type	Description
Report	Run the report, but with the opportunity to edit its current parameters (if present). If you choose this option, the Report Parameters dialog is displayed before the report is run. You can override the default report parameters for just this run of the report.
Report with defaults	Run the report directly, using its defined parameters, if present. For focused reports, this is the only option.
Delta report	For reports based on bar charts, run the report after selecting another report as the comparison for the delta.

3. In the Report Parameters dialog, enter settings in Common Parameters and Custom Parameters sections as appropriate.



Note: Use the dialog to override any predefined parameters for the report, but only for this particular run.

Report Parameter Values

Common Parameters	Description
Report Format	From the drop-down menu, select one of the following report output formats: • pdf - Outputs the report as an Adobe PDF file. • xls - Generates a Microsoft Excel file for tables and charts. Note: XLS reports you run with <i>Microsoft Excel 2002</i> might have page break format problems (misalignments, column spillover) due to default page size settings in Excel. To correct this problem, open the resulting XLS report in Excel, choose File > Page Setup from the menus, change the paper size to Letter (instead of Legal), and click OK to save your changes. The report has the appropriate page break formatting. <i>This problem does not occur in newer versions of Microsoft Excel.</i> Note: XLS report formats display speedometer charts as pie charts. This is a known limitation in Microsoft Excel. • rtf - Produces a rich-text format document. • csv - Creates tabular data as a list of comma-separated values. Note: Reports generated in CSV format are not the full equivalent of exports to other formats like PDF or HTML. CSV format is useful for loading report data into a spreadsheet for further manipulation. Since CSV is meant to contain tabular data, only the table data of a report is normally useful. Therefore, ArcSight exports only the table data portion of a report to CSV format, ignoring any other report information such as charts or text, including report titles. • html - Generates the report in a Web page displayed by the default web browser. Note: Running a report with a huge row limit (10,000 rows) in HTML format may cause the browser to be unresponsive. In such cases, run the report in PDF. Your selection affects your choice for e-mail formats. See description for Email Format, another report parameter, in this table. The output file is created at report runtime and is stored in the Archived Report Group corresponding to the report's group. See the Archives tab on the report's edit panel.
Page Size	From the drop-down menu, select a paper size.
Run as User	This option is not available for case reports. Run the report as a particular user. From the drop-down menu, select the user name by which you would like to run the report. For example, this option would allow an administrator for a Managed Security Service Provider (MSSP) to run report for a customer. The administrator would need write permissions to the user.

Report Parameter Values, continued

Common Parameters	Description
Email to	You can have the report sent as e-mail to one or more ArcSight users. From the drop-down menu, select the Console users to whom the report should be e-mailed. The selection list is read from the Users resource). The recipient will only see his or her user name in the To field even if there are multiple recipients for this report. Handling empty reports: If you are emailing reports, empty reports will also be sent. This is determined by the server property, <code>report.scheduler.notify_empty_reports</code>, which is set to true. If you do not want empty reports to be sent, add the property to the <code>server.properties</code> file and change the setting to false. Follow the instructions in the <i>ESM Administrator's Guide</i> on how to edit this file. The details are in the guide's topic, "Managing and Changing Properties File Settings."
Email addresses	Send the report to one or more comma-separated or semicolon-separated e-mail addresses. This option does not require the recipient to be an ArcSight user. Note: The recipient will only see his or her e-mail address in the To field even if there are multiple recipients for this report.
Email Format	Specify how the report is to be accessed by the recipient. • Select Send URL if you want to point users to the report. Use this option if the report is large and is saved (archived) to a network-accessible location You can provide URLs for all report formats: PDF, XLS, RTF, CSV, and HTML. • Select Attach Report if you want to send the report directly to the user's e-mail box. You can only attach PDF, XLS, RTF, and CSV report formats. • Select Attach Compressed Report if you want the PDF, XLS, RTF, or CSV report to be compressed (zipped) first before mailing. Note: The notification automatically uses the URL if you select an email format for an unsupported report format, or if the attachment is too large to be sent. You can adjust the maximum attachment size up to 100 MB with the property, <code>report.upload.maxFileSize</code>. Set this property to the size you want in the <code>server.properties</code> file. Refer to the <i>ESM Administrator's Guide</i> for procedures to edit properties files.
Email Subject	Specify the subject on the notification. Defaults to the report's Name attribute (denoted by <code>\$ReportName</code>). If you want to use a customized subject, type the text either in addition to the default or replace the default text entirely.

4. Click **Save Output** if you want to save a copy of the report to disk.

If you select this option, additional archive parameters are displayed. You can override any of these defaults also. You can:

- Change the group in which to archive the report.
- Provide a report name other than the default.

- Specify an expiration time at which to discard the report from the archive. By default, the report is saved in the archive for 6 months from the time it was run.



Tip: You can use Velocity template references for fields that accept text, such as Archive Report Name and Archive Report Expiration Time. See ["Velocity References for Reports" on page 1091](#) for details.

5. In the Report Parameters dialog box, enter new parameters if available and appropriate.
6. Click **OK**.
7. In the options dialog box click **Open** to open the report, **Save** to choose a location and format for the output file, or **Cancel** to quit. The Save option applies to all but HTML files.

Running a Delta Report

Purpose:

To show the difference between two sets of parameters of an existing report within a single comparative report. You can also set up a delta report to run and archive on a schedule. See also ["Archiving and Scheduling Reports" on page 459](#) and ["Scheduling Report Tasks" on page 461](#) for more information.

Requirements:

- Delta reports are available only for reports with a bar, 3D bar, or inverted bar chart.
- The report must contain one chart only (no tables).
- The X and Y axis must have at least one column each, and no Z-axis.
- The chart must not have any summary function or top N filter applied.

For more information about creating reports with these characteristics, see the ["Binding Data to the Report" on page 398](#).

Where: Navigator > Resources > Reports

1. Right-click a report and select **Run**, then **Delta Report**.
2. Select the parameters for the first report, select a report format from the drop-down menu, and click **OK**.
3. Select the parameters for the second report and click **OK**.
4. Select **Save Output** if you want to save a copy of the report to disk.

If this option is selected, additional archive parameters are displayed. You can override any of these defaults also. You can select a group in which to archive the report, provide a

report name, and specify an expiration time at which to discard the report from the archive. By default, the report is saved in the archive for 6 months from the time it was run.



Tip: You can use Velocity template references for fields that accept text, such as Archive Report Name and Archive Report Expiration Time. See ["Velocity References for Reports" on page 1091](#) for details.

The Report Viewer appears and displays the delta report. The report shows the difference between two sets of parameters used on a single report. The report also shows the data for each of the parameters.

When a delta report is run or archived, an internal event is sent to the ArcSight Manager. This event contains the following data fields and values:

Delta Report Event-data Field	Description
Event Name	Delta Report Generated (Report: <ReportName>), where <ReportName> is the name of the report.

Running Reports from a Grid View

You can define reports on demand based on specific events in grid views in the Viewer panel:

- ["Running a Rule Context Report" below](#)
- ["Running an Event Context Report" on the next page](#)
- ["Running a Channel Report" on the next page](#)
- ["Running a Query Viewer Report" on page 453](#)

Running a Rule Context Report

ESM generates a correlation event when a rule finds matching base events and triggers an action.

Purpose:

To generate a report containing the selected correlation event and associated base events that caused the rule to trigger.

Procedure:

1. In a grid view, select a correlation event marked with a **Flash** icon ().
2. Right-click it and select **Report > Rule Context Report**.
3. In the **Report Parameters** popup, enter the time, in minutes, before and after this event's occurrence and click **OK**.
4. You can choose to **Open** or **Save** the report file.

A report showing the correlation event and the events that triggered the rule appears.

Running an Event Context Report

Purpose:

To generate a report that includes the events that occurred within a specified time before and after the event you select in the grid.

Procedure:

1. In a grid view, select an event.
2. Right-click and choose **Report > Event Context Report**.
3. In the **Report Parameters** popup, enter the time, in minutes, before and after this event's occurrence and click **OK**.
4. You can choose to **Open** or **Save** the report file.

Running a Channel Report

Purpose:

To generate a report that includes **all** events in the channel, not just the selected event.

Procedure:

1. In a grid view, select an event.
2. Right-click and select **Report > Channel Report**.
The Report Parameters popup is displayed, and its fields are automatically populated with the event data fields. You can enter new parameters to limit or extend the report.
3. Select a **Report File Format** from the drop-down menu.

4. Click **OK**.
5. You can choose to **Open** or **Save** the report file.

Running a Query Viewer Report

Purpose:

To generate a simple report based on query viewer results, therefore providing a convenient way to share the result data.

Query viewer reports are limited to displaying data from the single query covered by the query viewer and retain the format of the chart or table in which the query viewer results are displayed.



Tip: On row limits

The report display format is based on the display chosen for the query viewer result. For example, if you chose to view query data as a *pie chart*, the generated report shows the same pie chart view. To generate a report showing results for the same query as a *bar chart* or *table*, you must then re-run the query viewer (<Query Viewer> > View Data as) in one of those formats, and then generate the report from that view.

The report contents might not include as much data as the query viewer result shown in the Console for these reasons:

- Reports on pie charts and bar charts have a default row limit of 25 up to 99 rows as the maximum. This is user-configurable. You can set a higher or lower row limit on the Report Parameters dialog you get when you run the report. (See the procedure below.)
- Reports on *tables* can accommodate up to 10,000 rows.

To generate a report on a query viewer:

1. Right-click the query viewer results table or chart (anywhere in the Viewer panel) and select **Report**.
2. Specify the options on the Report Parameters dialog or take the defaults and click **OK**.

Report Parameter Values

Common Parameters	Description
Report Format	From the drop-down menu, select one of the following report output formats: pdf - Outputs the report as an Adobe PDF file. xls - Generates a Microsoft Excel file for tables and charts. Note: XLS reports you run with <i>Microsoft Excel 2002</i> might have page break format problems (misalignments, column spillover) due to default page size settings in Excel. To correct this problem, open the resulting XLS report in Excel, choose File > Page Setup from the menus, change the paper size to Letter (instead of Legal), and click OK to save your changes. The report has the appropriate page break formatting. <i>This problem does not occur in newer versions of Microsoft Excel.</i> Note: XLS report formats display speedometer charts as pie charts. This is a known limitation in Microsoft Excel. rtf - Produces a rich-text format document. csv - Creates tabular data as a list of comma-separated values. Note: Reports generated in CSV format are not the full equivalent of exports to other formats like PDF or HTML. CSV format is useful for loading report data into a spreadsheet for further manipulation. Since CSV is meant to contain tabular data, only the table data of a report is normally useful. Therefore, ArcSight exports only the table data portion of a report to CSV format, ignoring any other report information such as charts or text, including report titles. html - Generates the report in a Web page displayed by the default web browser. Note: Running a report with a huge row limit (10,000 rows) in HTML format may cause the browser to be unresponsive. In such cases, run the report in PDF. Your selection affects your choice for e-mail formats. See description for Email Format, another report parameter, in this table. The output file is created at report runtime and is stored in the Archived Report Group corresponding to the report's group. See the Archives tab on the report's edit panel.
Page Size	From the drop-down menu, select a paper size.
Run as User	This option is not available for case reports. Run the report as a particular user. From the drop-down menu, select the user name by which you would like to run the report. For example, this option would allow an administrator for a Managed Security Service Provider (MSSP) to run report for a customer. The administrator would need write permissions to the user.

Report Parameter Values, continued

Common Parameters	Description
Email to	You can have the report sent as e-mail to one or more ArcSight users. From the drop-down menu, select the Console users to whom the report should be e-mailed. The selection list is read from the Users resource). The recipient will only see his or her user name in the To field even if there are multiple recipients for this report. Handling empty reports: If you are emailing reports, empty reports will also be sent. This is determined by the server property, <code>report.scheduler.notify_empty_reports</code>, which is set to true. If you do not want empty reports to be sent, add the property to the <code>server.properties</code> file and change the setting to false. Follow the instructions in the <i>ESM Administrator's Guide</i> on how to edit this file. The details are in the guide's topic, "Managing and Changing Properties File Settings."
Email addresses	Send the report to one or more comma-separated or semicolon-separated e-mail addresses. This option does not require the recipient to be an ArcSight user. Note: The recipient will only see his or her e-mail address in the To field even if there are multiple recipients for this report.
Email Format	Specify how the report is to be accessed by the recipient. • Select Send URL if you want to point users to the report. Use this option if the report is large and is saved (archived) to a network-accessible location You can provide URLs for all report formats: PDF, XLS, RTF, CSV, and HTML. • Select Attach Report if you want to send the report directly to the user's e-mail box. You can only attach PDF, XLS, RTF, and CSV report formats. • Select Attach Compressed Report if you want the PDF, XLS, RTF, or CSV report to be compressed (zipped) first before mailing. Note: The notification automatically uses the URL if you select an email format for an unsupported report format, or if the attachment is too large to be sent. You can adjust the maximum attachment size up to 100 MB with the property, <code>report.upload.maxFileSize</code>. Set this property to the size you want in the <code>server.properties</code> file. Refer to the <i>ESM Administrator's Guide</i> for procedures to edit properties files.
Email Subject	Specify the subject on the notification. Defaults to the report's Name attribute (denoted by <code>\$ReportName</code>). If you want to use a customized subject, type the text either in addition to the default or replace the default text entirely.

 **Tip:** If you click **Save Output** on the Report Parameters dialog, you get additional options for setting the archived report under the Save Output Parameters section.

When the report is ready, a dialog gives you the option of opening it to view it now or saving it to a location you specify through a file browser.

3. Select **Open** to view the report or **Save** to save it in a specified location.

Running Large or Complex Reports

A very large report (for example, a 500 MB PDF report) might require so much virtual machine (VM) memory that it can cause the ArcSight Manager to crash and re-start. To prevent this scenario, you can set up the Manager to expose a special report parameter for generating the report in a separate process. The separate process has its own VM and heap, so the report is more likely to finish. Even if the allocated memory is still not enough, the report failure will not crash the Manager. This option must be set up on the Manager to expose it in the ArcSight Console report parameters list.

The default server property is `report.canarchiveportinseparateprocess=false`. You need to change this to **true**.

The steps are as follows:

1. Refer to [ESM Administrator's Guide](#). In the Configuration chapter, follow the instructions in "Managing and Changing Properties File Settings" to add the property. Make sure to restart the Manager after making changes.

After setting the property to **true**, you will now have the ability to set the report parameter on the ArcSight Console.

2. On the ArcSight Console, open the report that you want to run in a separate process in the Report **Editor**, and click the **Parameters** tab. Set the parameter **Generate Report In Separate Process** to **true**.
3. Run the report. The report should run like a normal report, but it does not consume the resources of the Manager VM. See notes below for more information.



Tip: Here are more tips about the property described in this topic.

- If a report is saved with the parameter set to **true**, the report is archived as a separate process even if the property `report.canarchiveinseparateprocess` in `server.properties` is set back to **false** later on.
- This property indicates whether reports are allowed to be archived in a separate process. When this property is set to **true**, the option to run and archive the report in a separate process is available in the common properties in the Report Editor. Setting the value to **true** causes the report to be archived in a separate process. The benefit of archiving a report in a separate process is to avoid consuming Manager resources and potentially crashing the Manager.
- Use this parameter only in special circumstances as needed. For example, if archiving a report is causing the Manager to crash then you might apply this solution. Generally, if a report contains tables that have more than 500,000 rows with 4 or 5 columns per row it is likely that the report is large enough over-tax the Manager VM memory. However, the threshold may vary depending on the Manager heap size and the details and data in the tables so it is best to only resort to this solution if you encounter problems archiving a particular report.

Moving and Copying Reports

Purpose:

To move or duplicate report definitions to better organize your work, to publish your definitions, or to make editable copies of enterprise reports.

Where: Navigator > Resources > Reports

Procedure:

1. Drag and drop the report into another group.
2. Select **Move** to move the report, **Copy** to make a separate copy of the report, or **Link** to create a copy of the report that is linked to the original report.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.


Importing and Exporting Reports

To import and export reports across compatible ESM systems, use the Packages feature. See ["Managing Packages" on page 730](#) for details.

Managing Report Groups

Report groups store similar reports, and control access to reports, using access control lists (ACLs). When editing access control permissions, permissions given to a report group are also given to all groups and reports within that group.

Groups and reports can be managed with drag and drop functionality. You can move or copy groups and reports into other groups from the Reports resource tree. If a group is deleted, the reports within that group are also deleted.



Where: Navigator > Resources > Reports

To create a report group:

1. Right-click a Reports group and select **New Group**.
2. Enter a report group name in the text field.
3. Press **Enter**.

To rename a report group:

1. Right-click a Reports group and select **Rename**.
2. In the "name" text field, rename the group.
3. Press **Enter**.

To edit a report group:

1. Right-click a Reports group and select **Edit Group**.
2. In the Report Editor, edit the **Name** and **Description** text field.
3. Click **OK**.

To move or copy a report group:



Note: To copy multiple resources at once, use Copy and Paste. You can drag and drop only one resource at a time.

1. Drag and drop the Reports group into another Reports group.
2. Select **Move** to move the group, **Copy** to make a separate copy of the group, or **Link** to create a copy of the group that is linked to the original group.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete a report group:

1. Right-click a Reports group and select **Delete Group**.
2. Click **Yes**.

Archiving and Scheduling Reports

You can schedule reports to archive automatically using the scheduler. The scheduler accepts multiple schedules by year, month, week, day, or hour. For example, a report can be archived automatically on the first of January at both 5 AM and 6 PM. The scheduler sends e-mail notifications informing users when a scheduled report has been archived.

Archiving a Report



Note: If you do not want to archive empty reports:

By default, the scheduler archives reports even if they are empty. If you do not want this behavior, enter this property in `server.properties`:

```
report.scheduler.archive_empty_reports=false
```

Refer to "Managing and Changing Properties File Settings" in the [ESM Administrator's Guide](#) for instructions on how to change settings properly.

Where: Navigator > Resources > Reports

To archive a report:

1. Right-click the report and select **Schedule for Archiving > Report**.

The report editor opens at the **Jobs** tab.

2. Click **Add** and select either **Schedule Report** or **Schedule Delta Report**.



Tip: The option to **Schedule a Delta Report** job is available only for certain types of event-based reports, and only when a previously-run report is available in the archives. Otherwise, clicking **Add** on the Jobs tab takes you directly to the job scheduler to schedule a standard report. For more information about delta reports, see ["Running a Delta Report" on page 450](#).

3. Enter a name and description for the job.
4. In the Jobs scheduler, click the link labeled **Click here to set up schedule frequency** to open the Job Frequency dialog, and configure the schedule.

In the Job Parameters section, select or enter values for the parameter fields as necessary. For date parameters, enter values in the text fields, click the drop-down arrows or click the time buttons to select a time range. For time data, you can enter a specific value, such as 8:54:00 AM or you can use special timestamp variables.

Click **OK** to save changes to the Job schedule.

To view all scheduled jobs, click the **Open scheduled jobs list** tool button (). The scheduled tasks are listed in the Viewer panel under “Current Jobs.”

For more information on setting up and viewing scheduled jobs, see ["Scheduling Report Tasks" on the next page](#).

5. Back in the report editor Jobs tab, under the **Job Parameters** section, enter values for the report parameters by clearing the Use Default check marks, or keep the defaults. You can set the report format, e-mail options, output parameters, start and end times, and so on. These are the same parameters described in ["Setting Default and Custom Report Parameters" on page 415](#).

If you chose any parameterized condition while creating this report, that condition is displayed at the top of the tab. An example of a parameterized condition is: detect time between \$CurrentDate-1d and \$CurrentDate. If such parameters exist, they are used for both immediate as well as scheduled generation of reports. While scheduling reports for archiving, these parameters are displayed in the Edit Parameters UI. It is possible to independently modify the dates specified in the parameter text fields. In addition to relative dates, you can specify absolute dates as parameter values. Examples of valid absolute dates are: 01/01/2015 and 01/01/2016 11:00:00 AM.

6. Click **Apply** or **OK** on the report Editor to save your changes for this report.

Displaying an Archived Report

Where: Navigator > Resources > Reports > Archives tab

Right-click an archived report and select **Show Archive Report**.

Scheduling Report Tasks

You can schedule some tasks to occur automatically. Specifically, this feature is available for archiving reports individually or by group, for taking Threat Detector snapshots, and for scheduling rules. This topic discusses the scheduler as it relates to scheduling reports .

Topics include:

- ["Scheduling Individual-Report Archiving" on the next page](#)
- ["Scheduling Report Archiving by Resource Group" on page 463](#)
- ["Standard Time Transitions" on page 464](#)

For more information on job scheduler in general, see also ["Job Scheduler" on page 1003](#).

Scheduling Individual-Report Archiving

Where: Navigator > Resources > Reports

To schedule archiving of individual reports:

1. Right-click the report you want to schedule.
2. Select **Schedule for archiving**, then **Report** or **Delta Report** for delta reports.
The report editor opens at the Jobs tab.
3. Click **Add**.
4. Enter a name and description for the job.
5. In the Job Parameters section, select or enter values for the parameter fields as necessary.
6. In the Summary section, click the link labeled **Click here to set up schedule frequency**.
7. In the Job Scheduler dialog, select the desired frequency and enter the associated settings:

Option	Description
Schedule Frequency	Choose a timing for the schedule. The typical choices are self-explanatory: Hourly, Daily, Weekly, Monthly, and Yearly. For each timing, enter additional settings such as occurrence (Every) and time. For Monthly, also specify the day of the month. For Yearly, specify the month and date. Caution: If you are scheduling a yearly report on a leap year and you choose February 29 as the month and date, the Console will not save your setting. To ensure that yearly reports are covered on leap years, choose a different date; or choose either February 28 or March 1 if you want your schedule to be as close as possible to February 29.
Schedule Range	Select the start and end date for the report run.

8. Repeat as required to add another schedule for the same group.
9. Click **OK**.
Reports can be archived in PDF, HTML, Excel, Comma Separated Value (csv), or Rich Text Format (rtf). The default PDF format should be used when archiving reports. Compared to PDF reports, other reports may lose formatting information and appear differently. In addition, Excel format is more memory-intensive than PDF.
10. Select the **e-mail scheduled reports to** check box and a user from the drop-down menu to automatically send an e-mail notification when the report is archived.
The user receives an e-mail notification stating that the report has been successfully archived. The e-mail also contains a URL to the report so that the user can view the report

from the URL. The e-mail notification is sent to the e-mail address listed in the user's profile. The user must have an e-mail address in their user profile.

11. For the **Archive Folder** text field, click the archive report group button to select where to list the archived report.
12. In the Archive Report Selector, select a report archive group and click **OK**.
13. In the Report Parameters window, click **Update**.
14. For delta reports, in the Schedule Summary, right-click **Default** under the **Param Set 2** column and select **Edit Parameters** to change the second parameter set, if any. Click **Update**.
15. In the Schedule Summary, click **Close**.



Tip: You can use Velocity template references for fields that accept text, such as Archive Folder and Archive Report Selector. See ["Velocity References for Reports" on page 1091](#) for details.

Scheduling Report Archiving by Resource Group

Where: Navigator > Resources > Reports

1. Right-click the report group and select **Schedule for archiving>Report group**.
2. On the **Jobs** tab, click **Add**.
3. Enter a name and description for the job.
4. In the Job Parameters section, select or enter values for the parameter fields as necessary.
5. In the Jobs scheduler, click the link labeled **Click here to set up schedule frequency**.
6. In the Job Scheduler dialog, select the desired frequency and enter the associated settings:

Option	Description
Schedule Frequency	Choose a timing for the schedule. The typical choices are self-explanatory: Hourly, Daily, Weekly, Monthly, and Yearly. For each timing, enter additional settings such as occurrence (Every) and time. For Monthly, also specify the day of the month. For Yearly, specify the month and date. Caution: If you are scheduling a yearly report on a leap year and you choose February 29 as the month and date, the Console will not save your setting. To ensure that yearly reports are covered on leap years, choose a different date; or choose either February 28 or March 1 if you want your schedule to be as close as possible to February 29.
Schedule Range	Select the start and end date for the report run.

7. Click **OK**.
8. Repeat as required to add another schedule for the same group.

Standard Time Transitions

If the trigger time for a particular scheduled task run happens to fall during the transition time between daylight saving time (DST) and standard time (ST), the interval for that run will be different than expected.

Time zones that honor DST have a period of time that occurs twice during the transition from DST to ST. For example, in the US when changing from DST to ST, this hour occurs once while the DST is still in effect and again after switching to the Standard Time. The transition period occurs at 2 a.m., therefore 1:00:00 a.m. - 1:59:59 a.m. occurs twice (1:00:00 a.m. PDT - 1:59:59 a.m. PDT and 1:00:00 a.m. PST - 1:59:59 a.m. PST), where 1:00 a.m. PST is 60 minutes after 1:00 a.m. PDT. In this example, if the scheduled task is due to trigger any time between 1:00:00 a.m. - 1:59:59 a.m., the interval for that particular run of the scheduled task will not be as expected.

Similarly, when the time changes from ST to DST, the 1:00:00 a.m. - 1:59:59 a.m. hour does not occur at all. The local time changes directly from 1:00 a.m. to 2:00 a.m. So, if your scheduled task run was scheduled to trigger between 1:00:00 a.m. - 1:59:59 a.m., the interval for that particular run will be off by an hour.

The interval calculation for subsequent scheduled runs do not get affected.

Currently, there are four time zones that are not supported in ESM:

- Kwajalein
- Pacific/Kwajalein
- Pacific/Enderbury
- Pacific/Kiribati

These time zones fall in two countries, Marshall Islands and Kiribati.

Viewing an Archived Report

Where: Navigator > Resources > Reports > Archives tab

To view an archived report:

1. Navigate the Archived Reports tree to find the archived report you want.
2. Right-click the report and select **Show Archive Report**.

The report is displayed in the Viewer.

If you do not find the report you are looking for, you might want to check to see if it has run yet.

To view all scheduled jobs:

Click the **Open scheduled jobs list** tool button (🔍). The scheduled tasks are listed in the Viewer panel under Current Jobs.

Editing a Report Archiving Schedule

You can change the archiving schedule for report definitions in your Reports resource folders.

Where: Navigator > Resources > Reports

1. Right-click a report and select **Schedule for archiving**, then **Report** or **Delta Report** as applicable.
2. In the Schedule Summary, right-click in the braces { } column and select the **Parameters** option to change report parameters set for the specific scheduled report. To delete a current scheduled archive report, right-click in the braces { } column of an existing schedule and click **Delete**.
3. To change the interval scheduling of a report, click the report interval button and **Yearly**, **Monthly**, **Weekly**, **Daily**, or **Hourly**, click the date and time buttons.
4. If editing within the same time frame, click the **Month**, **Date**, **Day**, **Hour**, **Min**, **AM/PM** buttons to specify changes to the report schedule.
5. Click **OK**.

Editing Report Archiving Parameters

You can change the archiving parameters of the report definitions in your Reports resource folders.

Where: Navigator > Resources > Reports

1. Right-click a report and select **Schedule for archiving**, then **Report** or **Delta Report**.
2. Right-click in the braces { } column for a scheduled report and select the **Parameters** option.
3. In the Report Parameters window, enter in the report parameter text fields, if any.
For date and time data fields, such as Detect Time, you can type an actual date value, such as 10/12/2016 8:54:00 AM, or you can use special timestamp variables.

4. Select the **E-mail scheduled reports to** check box, and a user from the drop-down menu, to automatically send an e-mail notification when the report is generated.

The user receives an e-mail notification stating that the report has been successfully archived. The e-mail also contains a URL to the report so that the user can view the report from the URL.

The e-mail notification is sent to the e-mail address listed in the user's profile. The recipient must have an e-mail address in their user profile.

5. For the Archive Folder text field, click the archive report group button to select where to list the archived report.
6. In the Archive Report Selector, select a report archive group and click **OK**.
7. In the Report Parameters window, click **Update**.
8. For delta reports, in the Schedule Summary, right-click **Default** under the **Param Set 2** column and select **Edit Parameters** to change the second parameter set, if any. Click **Update**.
9. In the Schedule Summary, click **Close**.

Deleting a Report Archiving Schedule

Where: Navigator > Resources > Reports

You can remove individual archiving schedules for reports in the Scheduled Tasks list.

1. Right-click a scheduled report (showing a calendar icon), select **Schedule for archiving**, then **Report** or **Delta Report**.
2. On the line for the schedule to remove, right-click in the braces { } column and select **Delete**.
3. Click **Delete** in the dialog.

Chapter 15: Building Trends

A trend is an ArcSight resource that defines how and over what time period data are aggregated and evaluated for tendencies and patterns. A trend executes a specified query on a defined schedule and time duration.

The ArcSight trends engine evaluates source data for trends based on:

- Event conditions, such as
 - Number of worm outbreaks
 - Incident time-to-close
 - Number of cases closed
- Common network elements, such as
 - Operating system
 - Business role
 - Regulatory compliance relevance

Trends can be used as the primary data source for a report, or used as the data source input to another query which is then used in a report (perhaps along with other queries or trends).

See also:

- ["Building Reports" on page 388](#) for an overview of all reporting tasks and tools
- ["Understanding the Reporting Workflow" on page 388](#) to see how trends fit in to the process of creating a report

How Trends Work

A trend references a query, specifies a schedule on which the query automatically triggers, and provides mechanisms for efficiently storing, viewing, and leveraging the trend results for reporting. The trend results are stored in a trend table in the ArcSight database and can therefore be queried.

You can set trends to run indefinitely or to end at a specified date and time. A trend can start retrieving historical data from logs, start with current events, or at some specified time in the future. You can also specify advanced options on how and when to build tables and store data.

After trend data is collected, you can view the results in the Data Viewer table and generate a trend report that displays the results in tables and graphs.



Depending on the data gathered by the base query, the trend is either *snapshot trend* or an *interval trend*.

Snapshot Trend

A *snapshot trend* uses a query that operates on a fixed moment in time, for example, to gather information about assets on your network. Snapshot trends are built from queries based on assets, cases, active lists, notifications, or actors. For example, snapshot queries and the trends built from them would be used to determine metrics such as current number of assets, number of systems with a particular operating system, or number of systems with particular vulnerabilities. A snapshot trend operates on data in the current moment in time, and only collects data going forward. Therefore, you cannot use trends to answer the question, “how many assets were in this zone a month ago?” You can use trends to collect data from this point forward, however, and in a month from now, you have a month's worth of data telling you how many assets were in this zone at regular intervals over the last month.

Interval Trend


An *interval trend* uses a query that operates on events that happen over a specified time window. For example, the trend gathers information about how many events of a particular description occurred daily over a 6-month period.

Interval trends are event based. For example, an interval trend using a base query with a time window could gather information to determine the number of login attempts in the past hour. You can refresh an interval trend manually as needed by selecting the trend in the Navigator and selecting **Refresh** on the right-click context menu.

Query-Trend Relationships in Reporting

A base trend is made of one query. Trends can be used as the primary data source for a report. Or, a trend based on one query can be used as the data source to another query that further refines the initial query result. A collection of trend queries (queries that use trends as their

data source) can provide focused views of a data set, which can then be fed into a single report or multiple reports.



For example, you could create a trend called “VPN Logins Outcome - Hourly” that references a query that returns all VPN login attempts, successful logins, and failed attempts. You could schedule the trend to run hourly. You can use this base trend directly in a report.

A more powerful approach would be to refine further the data results by creating three new trend queries, each of which takes the base trend as its data source, but then sets further conditions on the query data to return one specialized slice of the results. One query could return only login attempts, another only successful attempts, and another only failed attempts. You could then draw on four queries in a single or multiple reports to show different views of the data. (The base query would show all types of login events, and the other three would show the focused views.)

A single query or trend can feed data into multiple reports, and a single report can capture data from multiple queries and trends.

The ability to automate and refine queries by feeding them into trends and vice versa, along with the flexibility in populating reports, solves many typical enterprise security reporting challenges. You can build a trend that gets a daily event count, feed the trend into a query that sums up the daily counts to get a monthly event count, and even feed that monthly count query into another trend and so forth. Managed Security Service Providers (MSSP) can tier query-trend approaches to create focused reports for multiple customers built from what are initially broad range queries.

Managing Trends

Before you begin building a trend, make sure that you have a query defined that captures the data on which to build a trend. See ["Queries" on page 319](#) if you need more information.



Note: For a query used in a trend, the query and associated schema referenced in the trend are set at the time the trend was created. After the trend is created, you can modify some elements of the query if they do not affect the trend. For example, you can add or remove columns in the query if the related trend does not depend on them. Such modifications made to a referenced query are not reflected in the trend. If you modify aspects of the query that a trend depends on, the trend is disabled.

Creating or Editing a Trend

Where: Navigator > Resources > Reports > Trends tab



Caution: Do not exceed more than 10,000 resources in a group.

As a general rule, it is best to create new content under an individual folder.

The high-level steps for creating a trend are as follows:

1. To create a trend, right-click a trend group and select **New Trend**.
To edit a trend, right-click the trend and select **Edit Trend**.
2. In the Attributes tab, enter or change settings. See [Trend Attributes](#).
3. Verify that the trend schema represented by the selected Data Fields is appropriate.
4. Test the trend schema to make sure it is returning the expected data as described in [Testing a Trend](#).
5. Define or change the trend schedule as described in [Trend Schedule](#).
6. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
7. Click **Apply** or **OK**.



Caution: Do not click **Apply** or **OK** until you have defined the required values in the **Trend** section (trend name and query to use) and the trend schema in the Data Fields section of the Attributes tab. When you commit changes to the trend, the query and the schema are set and cannot be edited. If you decide to use a different base query or need to make a change to the schema, delete the trend and start with a new trend.

Defining Trend Settings

Use the Trend Editor to build a new trend or edit an existing one.

Refer to the following topics:

- ["Trend Attributes" below](#)
- ["Trend Schedule" on page 475](#)
- ["Trend Parameters" on page 476](#)
- ["Trend Actions \(Add to Active List\)" on page 477](#)
- ["Example: Populating Active Lists with Trend Results" on page 480](#)
- Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).


Trend Attributes

The following fields in the Trend section are required attributes to specify when creating a trend.

Trend Attributes

Field	Description
Name	Name for the trend. Spaces and special characters are allowed. The name you enter here is the alias that appears in data source pick lists in other editors.
Query	Specifies the query that this trend uses. See "Queries" on page 319 • Use the Query drop-down menu to select a base query as the source data for your trend. The arrow displays the Queries resource tree. • Alternatively, if the resource tree has too many subgroups to traverse and you know the query by name, see "Using the Advanced Selector While Editing Resources" on page 88. Trends tend to use up system resources, so make your queries for trends as refined as possible. Caution: • After the trend is created, you can add columns to the base query, but columns added to this query after the trend is created are not used by the trend. You can remove columns from the base query that are not used by the trend. However, you must create a new trend and select that modified query if you want to use the updated column definitions. • The Trend table has a limitation in that the table stores only a single element for Group variables (for example, GetGroupsOfAssets or FormatGroupsOfAssets). Therefore, if you are querying an active list containing a list field of resource references, the trend displays only one resource reference instead of the complete list. Query viewers and reports do not have this limitation, and are therefore a better choice for this scenario.
Enabled	By default, the Enabled check box is checked. This activates the trend to begin working on live data as soon as the trend is created. Uncheck this box if you want to experiment with the trend before pushing it live.

The following example shows a trend that uses the Events Count query as its basis.



The **Data Collection** section provides default values for row limit and query duration.

Data Collection Fields

Field	Description
Trend Interval	Time span over which the trend operates. The default is one hour. For example, if the query counts the number of logins, this setting counts the number of logins every hour.
Row Limit	Maximum number of rows of data the trend captures. The default number is 1000. You can modify the trend's row limit without affecting the base query's row limit.

The **Advanced** section provides optional settings to offset trend data collection and refresh trend data at a specified point in the future. By default, the offset and refresh values are set to **None**. You can keep the defaults or modify as needed.

Advanced Fields


Field	Description
Offset Collection By	Delays trend data collection by the time period specified. Offsetting trend data collection time enables you to compensate for events that arrive to the Manager late, either from a time zone lag or other data collection lag. Trend data collection starts after the time delay entered here. Enter a time delay and select Hours or Minutes from the drop-down menu. The default offset is None.
Refresh Data After	Triggers the system to re-evaluate the query data automatically at a later time to capture any additional events that may have come in late. Enter a refresh interval and select Hours or Minutes from the drop-down menu. The default refresh is None. Note: The Manager supports late arrival of events. For example, a SmartConnector can send a batch of events later if it is falling behind. Consider explicitly scheduling a refresh of trend data <i>only</i> if SmartConnectors frequently lag behind in sending events to the Manager. If SmartConnectors rarely go down and are generally on time delivering events, there is no need to set this option. Caution: Automatic data refreshes can be resource intensive. If possible, consider manually refreshing your trend for a special reason, such as, after system downtime that causes trends not to run on schedule. For manual refreshes, click the Refresh Trend Runs button on the trend editor's panel, then cancel when done.
Partition Size	Specifies the time range of partitions for this trend data, which in effect determines the partition size. The default “time slice” for trend tables is WEEKLY. That is, if the default setting is used, each partition would contain a week's worth of data. Partition size can be set to weekly or monthly. You can always modify the Partition size as needed by editing the trend definition. The Partition Size works in concert with the Partition Retention Period, described next.
Partition Retention Period (in days)	Specifies the number of days to retain the partitions from this trend as active in the ArcSight database. The default is 180 days. You can always modify the Partition Retention Period as needed by editing the trend definition. Note: The Partition Retention Period works in combination with the Partition Size. The system makes sure you always have as much data, if not more, than you specified in the configuration of these two settings. Similarly for factors such as time zones and daylight savings time, more data (never less) is retained. For example, if the Partition Size is set to MONTHLY and the Partition Retention Period is 45 days, the system stores two months' worth of data in two partitions. If the Partition Retention Period is set to 0 days, the data collected from one run of the trend is retained until the next partition is started. For example, if the Partition Size is MONTHLY and the Partition Retention Period is 0 days, then you keep one month's worth of data. Make sure that the trend start date is appropriate; a trend with a MONTHLY partition size, 0 days retention, and a start date near the end of the month would not maintain data for very long.

Advanced Fields, continued

Field	Description
Query Overlap Time	The query overlap time is the amount of time by which the next query should overlap with the previous query (overlapping the tail-end of the previous query). By default, queries for interval trends have no gaps and overlaps. The default overlap is 0 (None), which corresponds to the normal non-overlapping trend query case. By setting a query overlap time, you can configure a trend to support calculations like trailing <i>N</i>-day moving averages. The query overlap time extends the trend to include overlapping query ranges. For example, to collect moving average data over a 10 day period, you could run the query each day over the previous 10 days. A query overlap time set to 0 (the default) would result in non-overlapping runs, such that the query would run every 10th day over the previous 10 days. On the other hand, to get an overlapping trend run, you could specify a 9-day overlap. With this setting, the query would run every day (10 day query - 9 day overlap) over the previous 10 days. The trend would gather data every day for days 1-10, 2-11, 3-12, and so forth. Notes: - Do not use queries on the event table for anything longer than a day. Queries longer than a day should normally only run on other trend tables to allow the query to finish in a reasonable amount of time. - This option is enabled for snapshot trends.
Imported Trend Start Time	This applies to trends created and packaged here for use in another Manager. If the trend is packaged without schedule start and end times, the trend start time specified here is used when the trend is imported. Consider using a dynamic start time such as <code>\$Today-1w</code> to ensure that the trend does not use a very old start date, for example, a six-month old trend is yet to be imported. Otherwise, you run the risk of the trend not returning data because data no longer exists in the system. If the trend is exported without Schedule start and end times and no value is specified for Imported Trend Start Time, then when the trend is imported it defaults to use <code>\$CurrentDate</code> as the start time. With <code>\$CurrentDate</code>, the trend captures data starting from 12:00:00 AM of the current day. Note: The imported trend start time takes effect only if the trend is exported without Schedule start time. To exclude the Schedule start time from a trend upon export, set the package Format option to export. For information, see the description of the package's Format options in "Creating or Editing Packages" on page 731.
Imported Trend End Time	If the trend is packaged without schedule start and end times, the trend end time specified here is used when the trend is imported. If the trend is packaged without Schedule end time and no value is specified for Imported Trend End Time, then when the trend is imported it defaults to using no end time. With this setting, the trend runs indefinitely until it is manually disabled or edited to include an end time. Note: The imported trend end time takes effect only if the trend is exported without Schedule end time. To exclude the Schedule end time from a trend upon export, set the package Format option to export. For information on this, see the description of the package Format options in "Creating or Editing Packages" on page 731.

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

The **Data Fields** section is where you build the trend schema. This is populated automatically when you first select the query to use in this trend. The list shows the data fields collected by the query you chose. By default, all the query fields are selected for use in the trend. If you do not want to use a particular data field, clear (click to remove the check mark) the Use box for that item. Also, you can select which fields you also want to index. Indexing is done mostly for query efficiency, and TimeStamp is selected by default. It is helpful if the query you are using returns a large amount of data, and you want to run sub-queries on the data.



Name	Use	Index
TimeStamp	☒	☒
Attacker Address	☒	☐
Attacker Asset Name	☒	☐
Attacker Zone	☒	☐
Device Event Category	☒	☐
End Time	☒	☐

The **Summary** box at the bottom displays a summary of the query interval and the schedule on which it runs.

Trend Schedule

Click the trend editor's **Trend Schedule** tab to review or modify settings for the following parameters:

- **Schedule Frequency** - Specifies how often the query runs and gather data. The default is to run once every hour on the hour.
- **Schedule Range** - The Start field for the range specifies the oldest data to get. By default, the date and time the trend was created is used as the trend schedule start time. Make sure this date is within your storage's retention period; otherwise, the trend may not return any data from the oldest specified date.

By default, there is no End Date. This means that moving forward, the trend will continue to collect data.

With the default settings, this trend would collect data once every hour on the hour until it is disabled manually.

A Summary of the configured schedule is shown at the bottom of the tab.

Trend : ArcSight User Login Trends - Hourly

Attributes Schedule Parameters Actions Notes

Query Data Source: ArcSight User Logins - Last Hour
Trend Interval: 24 hours
Next Run: 19 Jun 2015 05:00:00 PDT

Schedule Frequency

☐ Hourly Every 1 day(s)
☒ Daily Occurs at 5:00:00

Schedule Range

Start: 8 May 2015 00:00:00 PDT
End: ☒ No End date
☐ End on 18 Jun 2018 10:25:42 PDT

Summary

Occurs Daily every 1 days(s) at 5:00:00 AM.
Schedule will start on 5/8/15 12:00:00 AM.

Refresh Trend Runs Test OK Cancel Apply Help

Trend Parameters

The trend editor's Parameters tab lets you further refine the query results in terms of row limits, time zone restraints, filters, and start and end times. If you set parameters in the base query used by this trend, those parameters appear on the Trend Parameters tab. In the trend, you can specify default parameters.

At Report building time, you can opt to run the report with the default parameters or “all parameters.” You can also further refine parameter details for a specific run of a report. For more information on specifying parameters in reports, see ["Setting Default and Custom Report Parameters" on page 415](#).



Note: Trend start/end times and row limits are used for gathering the data, and overwrite the start/end times and row limits set in the base query. If you do not customize the Trend Parameters, the defaults on this tab are used (not the start/end times and row limit on the Query General Attributes tab).

For reporting on the data after the trend has collected it, you can set new start/end times and row limit in the **Report Parameters** tab. The report parameters prescribe only the “outbound” or publishing data derived from the data already collected, not how the data is gathered.

Trend Actions (Add to Active List)

Trend actions give you the option to send specified columns (fields) in trend results to [Active Lists](#) (see "[List Authoring](#)" on page 486). You do this by defining an **Add to Active List** trend action. On the Actions tab for a trend, you can select to send data from one or more columns in the trend results to a specified active list.



Tip: Trend actions for active lists are similar to the **add to active list** rule action described in "[Rule Actions Reference](#)" on page 539. Unlike rules, however, **add to active list** is the only action available for trends, and the settings are not as fine-grained as for rules; for example, thresholds, number of events, time units, and so on do not apply to trend actions.

How Trend Actions are Useful (Summary Views and Rules)

The *Add to Active List* trend action provides a mechanism to get information from trends outside of, and in addition to, reports. Active lists updated by trends support summary views of information from multiple trends.

You can build a single active list that gets updates from multiple trends (each trend updating different columns in the active list). Also, a single active list can receive updates and show information from trends as well as from other sources (for example, rules). Alternatively, you can build multiple active lists that get updates from a single trend.

Perhaps most importantly, the ability to populate active lists with trend data makes trend results readily available for use in rules, filters, active channels, and so forth.

Example Use Cases

The following example use cases show how to leverage trend results in active lists:

- **Taking Action on Event-Based Trends.** An analyst wants to monitor the logins per hour by users based on their typical hourly login patterns. Flag anything that is above a threshold or more than n times a user's previous average.

The analyst sets up a trend to update the information in a trend table based on aggregation of per-user login events. The trend would have an **action** that updates an active list with the most recent results. Then, the analyst configures a rule to update another active list when a user logs on and another rule to compare the current login count against what is normal for that user. Any gross discrepancy can be used to trigger an alarm about a possible threat.

- **Taking Action on Asset-Based Trends.** An analyst wants to monitor assets by how vulnerable they are, and watch for "unusual activity" on especially vulnerable assets.

The analyst sets up a trend to check vulnerability counts on assets and log the top n most vulnerable assets on a daily basis. The active list would have an **action** to update an active list. Incoming events on assets would trigger rules that would check this active list against the particular device and, if present, trigger extra processing.

Plan and Define Active Lists with Fields Mapped to Trend

As a first step in setting up trend actions, determine which active lists the trend should populate and with what data. You might have existing active lists to add trend data to, or you could create new lists specifically for some trend results. (See ["Example: Populating Active Lists with Trend Results" on page 480](#) for an example of designing an active list based on the trend fields you want to monitor.)

Working with Trend Actions

Use the trend **Actions** tab to configure actions on a new or existing trend.

Requirement: An fields-based active list to which the trend adds entries

Where: Navigator > Resources > Reports > Trends tab


To define a trend action:

1. Right-click a trend and select **Edit Trend**.
2. In the Trend Editor, click the **Actions** tab.
3. Select the action **On Trend Run**, right-click, and select **Add to Active List**.




Note: You can only use a *fields-based* active list in a trend action (not event-based lists). For more information on types of active lists, see ["Creating or Editing an Active List" on page 486](#), especially the attribute [Data: Event-based, Fields-based](#).

4. Select the active list to be updated by the trend:



5. On the Add to Active List dialog, select fields from the trend on the right side to map to active list fields on the left.



What you are doing in this step is mapping trend column names to active list column names. All the “key” columns required by the active list must have trend columns mapped to them so that the active list entry (row) is correctly updated by the trend. However, not all of the active list value columns need to have trend columns mapped. Not specifying all the key columns is an error.

6. Click **OK** to add the action.

The action shows on the actions tab.

Note that you could add more actions here (by selecting the **On Trend Run** and clicking **Add**), edit this action, or remove it.

You can add multiple actions to a single trend (that is, configure a single trend to update particular columns in multiple active lists with trend results).

7. Click **OK**.

To edit a trend action:

1. Double-click the trend you want to edit.
2. Click the trend **Actions** tab.
3. Select the action you want to edit and click **Edit**.
4. On the Add to Active List dialog, make changes to the field mappings as needed and click **OK**.
5. Click **OK**.

To remove a trend action:

1. Double-click to the trend you want to edit.
2. Click the Trend **Actions** tab.
3. Select the action you want to remove and click **Remove**.
4. Click **OK**.


Example: Populating Active Lists with Trend Results

Suppose you want to monitor top failed user logins daily and send that data to an active list. (You could then configure rules to interact with the active list and trigger an alarm based on some threshold; for example, a single user with a certain number of failed logins per day.) To do this, you could create an active list with fields that map to a trend that monitors “top users with failed logins”.

To see the fields in this trend:

- Select the trend of interest.
- Select the trend, right-click and select **Data Viewer** from the context menu to display the trend results in the Viewer.

Your active list must have one or more of the trend's fields to capture relevant data in the list, as we'll show in the next section where we define the trend.



The screenshot shows the ArcSight Data Viewer interface. The title bar is 'Viewer'. Below it, there are two tabs: 'Top Users with Failed Logins Details' (selected) and 'Top Users with Failed Logins per Day Details'. The main content area displays the following information:

- Name:** Top Users with Failed Logins per Day
- Start Time:** 15 Jul 2015 00:00:00 PDT
- End Time:** 16 Jul 2015 11:39:17 PDT
- Filter:** No Filter
- 25 shown / 25 matches**

Below this information is a table with the following columns: TimeStamp, Day, Username, and Number of Failed Logins. The table contains 25 rows of data, sorted by the number of failed logins in descending order.

TimeStamp	Day	Username	Number of Failed Logins
15 Jul 2015 00:00:00 PDT	2015-07-15		55
15 Jul 2015 00:00:00 PDT	2015-07-15	lara	28
15 Jul 2015 00:00:00 PDT	2015-07-15	john	16
15 Jul 2015 00:00:00 PDT	2015-07-15	UNKNOWN	8
15 Jul 2015 00:00:00 PDT	2015-07-15	**unknown**	8
15 Jul 2015 00:00:00 PDT	2015-07-15	rajiv	8
15 Jul 2015 00:00:00 PDT	2015-07-15	shannon	7
15 Jul 2015 00:00:00 PDT	2015-07-15	michael	7
15 Jul 2015 00:00:00 PDT	2015-07-15	root	6
15 Jul 2015 00:00:00 PDT	2015-07-15	kashanmi	4
15 Jul 2015 00:00:00 PDT	2015-07-15	admin	3
15 Jul 2015 00:00:00 PDT	2015-07-15	dferrier	2
15 Jul 2015 00:00:00 PDT	2015-07-15	ram	2
15 Jul 2015 00:00:00 PDT	2015-07-15	oracle	2
15 Jul 2015 00:00:00 PDT	2015-07-15	sysadmin	2
15 Jul 2015 00:00:00 PDT	2015-07-15	fsaf	1
15 Jul 2015 00:00:00 PDT	2015-07-15	fdsfds	1

To continue with the example, we could create a fields-based active list with fields that map to the trend “Top Users with Failed Logins per Day” as follows.

Inspect/Edit

Event Inspector **Active List: Top Users with Failed Logins per Day AL**

Attributes **Notes** + Add Entry

Active List

- * Name: Top Users with Failed Logins per Day AL
- Optimize Data: ☐
- * Capacity (x1000): 10
- * TTL Days: 1
- * TTL Hours: 0
- * TTL Minutes: 0
- Allow multi-mappings: ☐
- Partially cached: ☐
- Time partitioned: ☐

Common

- Resource ID: HS7vNpDM6ABCD4zCcB+NbWA==
- External ID:
- Alias (Display Name):
- Description:
- Version ID:
- Content Version ID:

(Name)
(Description)

* Data: ☐ Event-based ☒ Fields-based ☒ Key Fields

Name	Type	Sub-type	Key-field
User Name	String		☒
Day and Time	Date		☐
Number of Failed L...	Long		☐

OK Cancel Apply Help

Name	Type	Key Field
User Name	String	This is the key field.
Day and Time	Date	
Number of Failed Logins	Long	

When the trend runs, it populates the active list with data on top users with failed logins by user name, and list the count of failed logins for each user along with date/time information. This active list could be used as the basis for rules, filters, active channels, and so on.

Viewer					
Top Users with Failed Logins Details					
Name: Top Users with Failed Logins			24 shown / 24 matches		
Last Update: 16 Jul 2015 11:33:13 PDT					
Filter: No Filter					
User Name	Day and Time	Number of Failed L...	Creation Time	Last Modified Time	Count
unknown	15 Jul 2015 00:00:...	8	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
admin	15 Jul 2015 00:00:...	3	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
dferrier	15 Jul 2015 00:00:...	2	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
dsfdsa	15 Jul 2015 00:00:...	1	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
fdgfdgd	15 Jul 2015 00:00:...	1	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
fdsfds	15 Jul 2015 00:00:...	1	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
fsaf	15 Jul 2015 00:00:...	1	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
fsdfds	15 Jul 2015 00:00:...	1	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
gfdgfd	15 Jul 2015 00:00:...	1	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
john	15 Jul 2015 00:00:...	16	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
kashanmi	15 Jul 2015 00:00:...	4	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
lara	15 Jul 2015 00:00:...	28	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
michael	15 Jul 2015 00:00:...	7	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
oracle	15 Jul 2015 00:00:...	2	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
rajiv	15 Jul 2015 00:00:...	8	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
ram	15 Jul 2015 00:00:...	2	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
ROOT	15 Jul 2015 00:00:...	1	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
root	15 Jul 2015 00:00:...	6	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
saf	15 Jul 2015 00:00:...	1	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
shannon	15 Jul 2015 00:00:...	7	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
Super	15 Jul 2015 00:00:...	1	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
supertest	15 Jul 2015 00:00:...	1	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
sysadmin	15 Jul 2015 00:00:...	2	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1
UNKNOWN	15 Jul 2015 00:00:...	8	16 Jul 2015 06:40:...	16 Jul 2015 06:40:...	1

Notes on Trend Action Behavior

- Trend updates either insert the row if new, or update an existing row. The update only populates / overrides the columns specified by the trend column mapping. Any other active list columns that do not have trend column mappings preserve their existing values. What this means is that it is possible for a single active list to be updated by multiple trends, each updating different columns. The active list is appropriately locked during read-modify-write cycle to avoid data corruption.
- A trend can be executed under a variety of circumstances, including refresh and backfill (data from the past). However, only the most recent data are entered into the active list. For example, no backfill data are added to the active list. A trend refresh run does not normally cause the active list to update, with the only exception being if it is the most recent data being refreshed.
- This trend action never removes entries from the active list. If the you want to have entries removed, use the active list's TTL (time-to-live) to have them expire. TTL setting is described in ["Creating or Editing an Active List" on page 486](#).

Testing a Trend

Purpose: To verify that the trend is returning the expected data properly.

Where: Navigator > Resources > Reports > Trends tab

1. Do one of the following:
 - Right-click the trend and choose **Test** from the context menu

Or

 - Click **Edit Trend** to display the Trend editor in the Inspect/Edit panel.
2. Click the **Test** button at the bottom of any of the trend editor tabs.

ESM evaluates the current event stream for matching events and populates the Test Trend pop-up dialog. The message “Success: x rows” at the bottom of the dialog tells you how many rows your trend returned.

The Test Trend sample shows a maximum of 25 rows. For interval queries, the sample also shows data from, at most, the last hour. If there is no match for the data, the trend returns no rows. This may mean that your current event query data contains no matching events or resources, or it may mean that your query needs to be refined.


Viewing Trend Data

Where: Navigator > Resources > Reports > Trends tab

1. Select the trend for which you want to view the data.
2. Right-click and select **Data Viewer** from the context menu.

This Trend Data Viewer opens in the Viewer panel and shows the query results. As with other ArcSight event viewers, you can select an event or group of events, right-mouse

click, and access various tools from the context menu to use for further investigation.



The screenshot shows the ArcSight Viewer window with the title bar 'Viewer'. The breadcrumb navigation is 'ArcSight User Status > Latest Events By Priority > VPN Logins Outcome - Hourly Details'. The main content area shows 'Name: VPN Logins Outcome - Hourly' and '871 shown / 871 matches'. Below this is a table with the following columns: Time, Number of Logins, Category Outcome, and Hour. The table contains 20 rows of data representing hourly login attempts and successes over a period from July 27 to July 28, 2016.

Time	Number of Logins	Category Outcome	Hour
27 Jul 2016 12:48:26 PDT	26	/Attempt	12
27 Jul 2016 12:48:26 PDT	100	/Attempt	13
27 Jul 2016 12:48:26 PDT	12	/Failure	12
27 Jul 2016 12:48:26 PDT	47	/Failure	13
27 Jul 2016 12:48:26 PDT	13	/Success	12
27 Jul 2016 12:48:26 PDT	54	/Success	13
28 Jul 2016 08:48:26 PDT	22	/Attempt	8
28 Jul 2016 08:48:26 PDT	80	/Attempt	9
28 Jul 2016 08:48:26 PDT	10	/Failure	8
28 Jul 2016 08:48:26 PDT	38	/Failure	9
28 Jul 2016 08:48:26 PDT	12	/Success	8
28 Jul 2016 08:48:26 PDT	42	/Success	9
28 Jul 2016 09:48:26 PDT	23	/Attempt	9
28 Jul 2016 09:48:26 PDT	78	/Attempt	10
28 Jul 2016 09:48:26 PDT	13	/Failure	9
28 Jul 2016 09:48:26 PDT	35	/Failure	10
28 Jul 2016 09:48:26 PDT	11	/Success	9
28 Jul 2016 09:48:26 PDT	42	/Success	10

Refreshing Trend Data

In addition to relying on the scheduled execution of a query according to its interval trend schedule, you can manually refresh the trend data at any time by using the trend refresh feature.

Purpose: Manually re-run a trend to compensate for events that arrive to the Manager late, either from a time zone lag or other data collection lag.

Where: Navigator > Resources > Reports > Trends tab

To refresh a trend table:

1. Do one of the following:
 - Click the **Refresh Trend Runs** button on the trend's Attributes tab for the selected trend.
 - In the Navigator, select a trend you want to refresh, right-click and select **Refresh trend runs....**

This opens the **Refresh Trend Runs** dialog which displays the trend query start times of the selected trend.
2. Select a period consisting of a start and end time stamp under **Show Trends From**, select one or more of the trend runs under **Choose Trend Runs**, and click **OK** to refresh the selected trend runs.

This executes the base query and refreshes the trend table on the selected runs.



Tip: You also can configure data collection to be offset by some time period to compensate for late arrival of events. For more information, see the Trend Attributes topic in ["Trend Attributes" on page 471](#).

Disabling or Enabling a Trend

Purpose:

After creating and testing trends in a development environment, you may want to package the trend and send the package to the production environment. If so, it is a good practice to disable the trend and then send in a package. After the package is imported and installed in the destination, the trend can be re-enabled so it can start collecting data.

You also disable trends for troubleshooting purposes, for example, it is using too much system resources and you want to investigate and re-configure the settings.

To disable an enabled trend:

Right-click the trend and select **Disable Trend**.

To enable a disabled trend:

First review the trend's start date on the trend editor's Schedule tab. By default, the start date is when the trend was created. If the start date is way in the past, for example over six months ago, note that the trend will try to backfill the data from that schedule. This will impact performance, and if the data is no longer available, you may not see any results. Before enabling an old trend, change to a start date that will realistically pull the data you need.

See the publication, *ESM Best Practices: Trends*, for a detailed discussion on backfilling.

Deleting a Trend

Before deleting a trend, make sure there is no query on this trend.

To delete a trend, right-click it and select **Delete Trend**.

If there is a query on this trend, you see a message that states "Dependent Resource Warning." If you want to continue, force the deletion by clicking **OK**.

Chapter 16: List Authoring

Active lists and session lists are important tools for tracking traffic with IP addresses of interest.

While you can manually update active lists, their real value comes when you author automatic, rule-driven lists with dynamic content.



Note: ESM stores list entries internally in the most efficient way possible, but the console does not reflect this order. Therefore, variable access to list entries (for example, `GetListElement`) does not necessarily match the order in which the console displays list entries.

Required Settings for Large Lists

By default, active lists and session lists each support 1 million entries. This section describes the required settings to support up to 5 million entries.

- You increase the list entries limit through the `activelist.max_capacity` and `sessionlist.max_capacity` property settings in the `server.properties` file, as in the following example:

```
activelist.max_capacity = 5000000
sessionlist.max_capacity = 5000000
```

- To have adequate memory for 5 million entries, use these settings in the `server.wrapper.conf` file:

```
wrapper.java.initmemory=32768
wrapper.java.maxmemory=32768
```

Both files are stored in the Manager's config directory. See "Managing and Changing Properties File Settings" in the [ESM Administrator's Guide](#).

- Set the Console's Java heap setting to **1536 MB**.

Creating or Editing an Active List

Purpose: Active lists are defined in conjunction with **rules** specifically tailored to interact with and populate the lists dynamically. See "[Using Rules to Populate an Active List](#)" on page 494.

Lists not driven by rules are empty or contain only manually added entries that have not timed out.

Where: Navigator > Resources > Lists > Active Lists tab

Procedure:

1. To create an active list, right-click an active list group and select **New Active List**.

To edit an active list, right-click an active list and select **Edit Active List**.

2. Set options as follows:

Active List Attributes

In this field...	...enter this
Name	Enter a name for the active list. This name identifies the active list in ArcSight list selector popups. Spaces and special characters are allowed.
Optimize Data	If you want to create a hash-based list, click Optimize Data to toggle it on. This option reduces the memory usage of an active list. It is useful for active lists with more than 1,000 entries or for lists that contain a large amount of information per entry. See "Optimize Data with Hash-Based Active Lists" on page 821 .
Capacity (x1000)	This setting indicates the maximum number of active list entries the system is to keep in memory. The default is 10,000. For most cases, 10,000 is appropriate, however, you may want to adjust this setting if the devices you are monitoring for this active list contain a lot of data to ensure you have adequate memory cache available. Notes: • This represents a limit on in-memory capacity only. If you also select Partially cached, more entries are retained but this has an impact on performance when it is necessary to retrieve active list items from the database. • If the maximum number of entries is reached, an existing entry is randomly selected and removed. For multi-mapped lists, removal is based on the key field; and starts when the number of keys exceeds capacity. • Capacity influences the maximum memory that can be consumed by the active list. The memory usage is proportional to the number of entries in the list, which usually are less than the capacity. Capacity affects memory usage, but has little if any impact on performance.
TTL Days, TTL Hours, TTL Minutes	TTL (Time To Live) means the items remain on the list for <i>at least</i> the amount of time you specify in Days , Hours , or Minutes . Use 0 (zero) to cause the field to never expire. The maximum number of days is 99999 .
Count Limit	Count Limit is used to limit the number of unnecessary updates to active list entries and improve performance. For example, if an On Every Event rule adds an entry to a list, but additional rules only check if an entry is in the list, not the count, there is no reason to update the count field of the entry every time. The Count Limit is a hard limit for the maximum count for an entry. A value of 0 (zero) indicates an unlimited count.

Active List Attributes, continued

In this field...	...enter this
Allow multi-mappings	Check this box to allow multiple instances of key pairings. This enables a single key, such as an actor attribute, to map to multiple values, such as a set of roles. You can use this to return a list of entries with the same value for the key field. For example, with multi-mappings enabled, you can create an active list that could return multiple roles for an actor named Clark Kent (reporter, superhero, space traveller) or multiple names associated with a farmhouse in Kansas (Clark Kent, Superman, Kal-El). Note: Don't use this setting if you are creating a Time partitioned active list.
Partially cached	When Partially cached is selected, additional entries beyond the in-memory Capacity (x1000) maximum are stored and retrieved from the database. Using partial caching increases overall capacity but can impact performance because it takes more time to retrieve list entries from the database. This setting is required by active lists that are Time partitioned. Note: There is a limitation when in-memory resources such as active channels and data monitors are used to return values from a partially-cached list. Only those values that are in the cache are returned. Reports and query viewers are not affected by this limitation because these resources query the database directly and do not use cache.
Time partitioned	A partially-cached, time-partitioned active list enables you to capture data over time. Without time partitioning, a partially-cached list requires constant retrievals from the database to update the entries, and flushing out old entries are done at random. With time partitioning, the cached data is segregated into partitions based on the list's timestamp (Date field) value. Time-partitioned list data are kept in memory, and older data are the first to age out of the list. This option requires that: • The list must <i>not</i> be multi-mapped. • Partially cached must be enabled. • The list must be fields-based (not event-based). Fields must include at least a date and a string field that are set as key fields. Without a date key field, the time partitioned setting is ignored.

Active List Attributes, continued

In this field...	...enter this
Case Sensitivity	You can optionally configure the list to be case-sensitive or -insensitive. Furthermore for case-insensitive lists, you can specify case-insensitivity for keys only, or for both keys and values. The feature enables you to store and look up values in lists regardless of case. Select one: • Case-Sensitive (the default) • Key Case-Insensitive • Key & Value Case-Insensitive Important: After you save the list, you cannot change this setting. If you want to revert the case sensitivity setting, define a new list instead. Cautions on case-insensitive active lists: • If your list is case insensitive, do not use the Optimize Data option. • Lookups on case-insensitive lists will slow down query and active channel performance. Make sure your queries and variables (used by channels) get values from <i>case-sensitive</i> lists.
Cache Model	The Cache Model determines how list data is accessed in a distributed ESM cluster. • When Read Optimized is selected, a local copy of the list data is held by each component accessing the lists. The local cache provides the best performance for rule filters and data monitors that reference the list. However, changes to a Read Optimized list require a short time to propagate to each local copy, so some events might be evaluated against stale list data. • When Write Synchronized is selected, a single cache is shared by all components, so any change to a list is simultaneously visible to all members of the cluster. However, accessing the list is slower. When using the Write Synchronized option, it is important that rule filters are structured in such a way that it minimizes access to the list. Filter clauses are usually evaluated in the order they appear in the user interface, so checking event field values such as Device Event Class ID before an <code>inActiveList</code> clause can reduce the overhead involved in accessing the active list.  Note: With the Write Synchronized cache model, changing the list capacity does not take effect until the cluster is restarted. The Cache Model setting does not apply to lists deployed to a stand-alone ESM system.

Active List Attributes, continued

In this field...	...enter this
Common and Assign fields	Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see "Common Resource Attribute Fields" on page 725 .
Data: Event-based, Fields-based	In the Data panel, choose Event-based or Fields-based lists. Your entries here determine what kinds of values your list is populated with. Caution: After you have selected your data fields and saved the active list, you cannot add, remove, or change existing data fields. • The Event-based option is convenient for choosing event attributes as found in existing events. When checking or adding to an event-based list, you only need to supply an event. This option is not supported in time-partitioned lists. • The Field-based option offers detailed event and attribute selection controls that involve mapping fields to field attributes. Use this setting for time-partitioned lists. Field-based lists that use "Key Fields" are known as active lists with values. (For more information, see "Active Lists with Values" on page 823.)

3. If list data is event-based:
 - a. Click **Select Fields**.
 - b. On the Field Selector panel, select one or more event fields for your list data collection then click **OK**. Then click **Apply** or **OK** on the Active List Editor panel to save your event-based list.
4. If list data is field-based:
 - a. Under the Name column, replace *<Enter Name>* with a descriptive name for the field. For a list of restricted characters, see ["Field Naming Restrictions" on page 506](#).
 - b. Select the data type and corresponding subtype as applicable:

Active List Column Types and Subtypes

Type	Subtype
Date	This field is required for time-partitioned active lists. Additionally, you must set this as a key field. If the time-partitioned list has no date or time-based field, time partitioning does not occur. This Date field is used as a default Timestamp value for interval-type queries on active lists.
IP Address	This field supports IPv4 or IPv6 address. If the value is an IPv6 address, the resulting address will be simplified if applicable. For example, 2001:db8:0000:0000:0000 will be displayed as 2001:db8::

Active List Column Types and Subtypes, continued

Type	Subtype
Double, Integer, or Long	Optionally select one of the numeric subtypes to accumulate values when the field is updated, for example, by a rule. If you do not select a cumulative numeric subtype, the entries are replaced when the list is updated. • SUM adds the existing value and the inserted value. For example, if type is Double, subtype is SUM, and the current value is 100.00, inserting a value of 50.0 results in a new value of 150.00. • MAX takes the greater of the existing and inserted values. For example, if type is Double, subtype is MAX, and the current value is 100.00, inserting a value of 50.0 does not change the current value because 100.00 is already a maximum of itself and 50.0. • MIN takes the lesser of the existing and inserted values. For example, if type is Double, subtype is MIN, and the current value is 100.00, inserting a value of 50.0 results in a value of 50.0. Notes: • The cumulative values feature is only available in fields-based active lists. • Do not use cumulative numeric fields as key fields. • If you are manually editing list entries for the cumulative numeric subtypes, the value you enter is the final value. This means accumulation of values does not occur with manual entry edits. • These cumulative numeric subtypes are not supported in multi-mapped active lists because new entry values for the same key add rather than modify the entries. • Trends <i>cannot</i> act on lists with cumulative numeric fields. • Rules and Threat Detector <i>can</i> act on lists that use cumulative numeric fields.
MAC Address	MAC address of the format consisting of six groups of two hexadecimal digits per group. Use hyphen (-) as separators. For example <pre>01-00-5E-90-10-FF</pre>
Resource Reference	Any ArcSight Resource such as asset, report, actor, and so on.
String	This is optional for lists in general but required, along with a Date field, if your list is time partitioned.

- c. Optionally check **Key Fields** to enable a per-field Key option, and then select one or more data fields that must be unique.

Important: Key fields *must* have values because key fields are used to uniquely identify a record.

For example, the ArcSight-provided active list ArcSight Foundation/Configuration Monitoring/Assets with Recent Configuration Modifications uses fields-based data, and keys on unique values for asset address, zone, and name.

Field-based lists that use **Key Fields** are known as active lists with values. (For more information, see ["Active Lists with Values" on page 823.](#))



Note: For key fields, here are best practices:

- For a time-partitioned active list, your key fields must be a Date field and a string field.
- Do not make cumulative numeric fields as key fields.

Database columns are defined after the list is created. After the new list is saved, you cannot add, remove, or change columns to the list.

5. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
6. Click **Apply** to save and continue editing or **OK** to save and close.

You can use the **Add Entries** button in the Active List Editor to manually insert values to the current active list. See ["Viewing and Editing Active List Entries" below](#).

Viewing and Editing Active List Entries

Where: Navigator > Resources > Lists > Active Lists tab

To add active list entries:

1. Right-click an active list and select **Edit Active List**.
2. On the Attributes tab, click **Add Entry**.
3. On the Active List Entry Editor panel, add the value for your data list and click **Add**.
4. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
5. Click **OK**.

To view active list entries:

Right-click an active list and select **Show Entries**.

The viewer panel displays the list's entries.

To edit active list entries:



Caution: If your fields-based active list contains numeric subtypes to accumulate values, be careful about manual changes. Your manually-entered value replaces the existing cumulative value, and your new value is not cumulative.

1. Right-click an active list and select **Show Entries**.
2. In the active list grid view, right-click an entry and choose **Edit**.

3. Click the entry's **Source Address** or **Count** to make changes.
4. Click **Modify** to change the existing entry, **Add** to post the changed entry as a new one, or **Delete** to remove an entry.

To refresh the active list view:

Active lists show results as of the time they opened for viewing, or the last time they were refreshed.

Click the **Refresh** button in the view header to update the contents.

To clear active list views:

While monitoring a particular active list view, you may want to see only traffic that happens after a certain point in time. You can accomplish this by clearing the view.

1. In the Navigator panel's Active List resource tree, select the active list to clear.
2. Right-click and choose **Clear Entries**.

To use the context menu on an active list view:

Below are right-click context commands available in active list views:

Menu Command	Description
New	Add an entry to the active list using the Active List Entry Editor.
Edit	Edit the selected entry using the Active List Entry Editor.
Delete	Remove the selected entry from the active list.

To use an active channel add to or subtract events from an active list:

You can add or remove event-attribute-based active list entries using selected events in the active channel. This feature automatically offers the name of the active list that is appropriate for the selected event.

See ["Adding Events from a Channel to an Active List" on page 497](#) for instructions.

To filter an active list:

In addition to the constraints of an active list itself, you can place a temporary filter on an active list view to aid your analysis. Such filters are not saved with the active list.

1. Open an active list in the Viewer panel as described above.
2. Click the **Filter** status description in the view header to open the Common Condition Editor. For example, the status **No Filter Defined**.
3. Use the Common Condition Editor as described in ["Creating or Editing a Filter" on page 306](#).

To customize active list columns:

You can modify active list views just like other grid views, as described in ["Customizing Columns" on page 250](#).

Using Rules to Populate an Active List

Purpose: To demonstrate how to populate active lists with the rule action. The example uses a rule that captures VPN login events and adds data to a list.

See ["Rule Actions Reference" on page 539](#) for more information on the Active List rule action.

The high level process involves:

1. Creating the active list that forms the *table* to store and display the data. The active list shows the number of logins by user name.
2. Creating the rule to capture VPN login events and send matching events to the list. The rules *populate* and update the list.

Example Active List

Purpose: To provide the table that stores login events captured by the rule.

Create the active list:

1. Create a fields-based active list named **VPN Events**.
2. Add fields named User Name and Category, both of type String.
3. Set User Name as the Key field:

* Data: ☐ Event-based ☒ Fields-based ☒ Key Fields			
Name	Type	Sub-type	Key-field
User Name	String		☒
Category	String		☐

Example Rule to Populate the Active List

Purpose: To create a rule that captures VPN login events and populate the VPN Events active list. Values found in Event Name and Category Device Group fields will be used indicators of such events. A matching event triggers the rule and populates the list as follows:


Populate this field in the Active List	With the value from this field in incoming events
User Name	Target User Name
Category	Category Device Group

Create the rule:

1. Create a standard rule, also named **VPN Events**.
2. Make your list case insensitive.

Set rule conditions:


Set **Conditions** to capture events when the **Event Name** contains Login and **Category Device Group** contains VPN. To capture event names from various sources that might be formatted differently (for example, in all upper case, all lower case, or initial capitalization), uncheck the Case-Sensitive (**Aa**) option next to the Event Name field. This shows up in the **Conditions Summary** tab as follows:



Tip: More fine-grained conditions logic (as used in this example) requires more processing and can have a performance impact. For example, using "<SomeField> **Contains** <SomeString>" for a field lookup requires more processing than writing a field lookup like "<SomeField> = <SomeString>".


Set rule aggregation:

1. On the **Aggregation** tab, select the fields for aggregation *only if they are identical*.
2. In the Add Field dialog, set Aggregation for event1 on Category Device Group and Device Custom String.



Set rule actions:

1. Activate the rule Actions trigger **On Every Event**. De-activate the other triggers.
2. Select **Add to Active List**.
3. Add values for **User Name** and **Category** to the active list. Map the fields as follows:
 - User Name: Device Custom String1
 - Category: Category Device Group



4. Click **OK** to save the rule.

Test the rule:

Drag-and-drop the rules into the Real-time Rules folder to deploy them. You are prompted to move, copy, or link the rule. Linking is often most efficient.

More details are in ["Deploying Real-time Rules" on page 555](#).

When the VPN Events rule is triggered, user names are added to the VPN Events active list:

User Name	Category	Creation Time	Last Modified Time	Count
FRED	/VPN	5 Dec 2015 14...	5 Dec 2015 14...	7
JASMINE	/VPN	5 Dec 2015 14...	5 Dec 2015 14...	7
JEFF	/VPN	5 Dec 2015 14...	5 Dec 2015 14...	9
JOHN	/VPN	5 Dec 2015 14...	5 Dec 2015 14...	2
LARA	/VPN	5 Dec 2015 14...	5 Dec 2015 14...	1
MICHAEL	/VPN	5 Dec 2015 14...	5 Dec 2015 14...	1

A logical next step in this example scenario would be to create another rule that checks if certain user names are showing up in the active list, and then takes some action (like sending an e-mail or adding those names to a "suspicious users" list, if appropriate).

Adding Events from a Channel to an Active List

From an event that shows up on an active channel's grid, you can select the option to add the event to an existing list or remove it, if that event is already an entry on an active list.

The default procedure is:

1. Open an active channel of events.
2. Right-click a specific event on the channel grid, and choose **Active List > Add To > Other**.

The Add to Active List dialog appears.



Note: The options Untrusted List and Trusted List apply to event fields that are address related.

3. Browse the Active Lists resource selector to locate your active list and select it.

The selected active list's entries are displayed .

4. Click **OK**.



Tip: If adding events from the active channel to existing lists is a frequent task for you, you can add your collection of frequently-used lists directly to the Active List option. That way, the lists are displayed instead of the active list resource tree.

Follow the instructions in ["Customizing the Default Selections for Active Lists" on page 54](#)

Moving or Copying an Active List

Where: **Navigator > Resources > Lists > Active Lists tab**

1. Drag and drop the active list into another group.
2. Select **Move** to move the active list, **Copy** to make a separate copy of the active list, or **Link** to create a copy of the active list that is linked to the original active list.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

Importing and Exporting an Active List

Where: Navigator > Resources > Lists > Active Lists tab

To import an active list:

You can import a comma-separated-value (CSV) file as data. This is useful if you have data from other systems that you want to import; you can use the import to populate your active lists.

1. Right-click an active list and select **Import CSV File**.
2. In the file browser, select the CSV file you want to import and click **Open**.

The Import Preview opens. If this is the file you want to import, click **OK** to add it to the active list.

3. Right-click the active list you just populated with the CSV file and select **Show Entries**. This displays the newly-added data from the CSV file in the Viewer panel as active list details.



Note: The default view limit is 2000 entries. To view more, specify the number of entries in your filter.

To export an active list:

In the active list viewer, you can export selected entries from an active list to a CSV file. This is useful if you want to manage active list data external to the ArcSight Console.

1. Right-click an active list and select **Show Entries**. The data in the active list is displayed in the Viewer panel as active list details.
2. On the active list detail in the Viewer panel, select one or more entries (rows of events).
3. Right-click and select either **Export CSV - Visible Columns** or **Export CSV - All Columns**.
4. In the file browser, go to the location where you want to save the exported data, enter a file name in the File Name field, and click **Save**. The entries you selected for export are saved as a CSV file.

Deleting an Active List

Before deleting an active list, make sure it is not being referenced by any rule action, trend action, or any other resource. Edit those resources if necessary.

1. Right-click an active list and select **Delete Active List**.
2. Click **Yes** to confirm.

Managing Active List Groups

Active list groups are created to store similar groups or active lists in a single location. Groups can be created within groups to meet enterprise needs.



Caution: Do not exceed more than 10,000 resources in a group.

Groups and active lists can be managed with drag and drop functionality. You can move or copy groups and active lists into other groups in the Active Lists resource tree. If a group is deleted, the active lists within that group are also deleted.



Note: To copy multiple resources at once, use Copy and Paste. You can drag and drop only one resource at a time.

Where: Navigator > Resources > Lists > Active Lists tab

To create an active list group:

1. Right-click a group and select **New Group**.
A name text field appears under the group you selected.
2. Enter a name.
3. Press **Enter**.

To edit an active list group:

1. To rename the active list group:
 - In the Active Lists resource tree, right-click a group and choose **Rename**, or
 - In the Group Editor, edit the **Name** field.
2. In the Group Editor, change the description as required.
3. Press **Enter**.

To move or copy active list groups:

1. In the Active Lists resource tree, navigate to a group and drag and drop it into another group.

2. Select **Move** to move the group, **Copy** to make a separate copy of the group, or **Link** to create a copy of the group that is linked to the original group.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete active list groups:

1. Right-click a group and select **Delete Group**.
2. Click **Yes**.

Managing Session Lists

While you can manually update session lists, their real value comes when you author automatic, rule-driven lists with dynamic content.

Related topics:

- ["Creating or Editing a Session List" below](#)
- ["Editing Session List Entries" on page 503](#)
- ["Understanding Session Correlation" on page 565](#)
- ["Example: Using Session Lists to Correlate Session Data on User Logins" on page 569.](#)

Creating or Editing a Session List

Purpose: Session lists are defined in conjunction with **rules** specifically tailored to interact with and populate the lists dynamically.

Where: **Navigator > Resources > Lists > Session Lists tab**

To create or edit a session list:

1. To create a session list, right-click a session list group and select **New Session List**.
To edit a session list, right-click the session list and select **Edit Session List**.
2. Set options as follows:

Session List Attributes

In this field...	...enter this
Name	Enter a name for the session list. This name identifies the session list in ArcSight pick lists. Spaces and special characters are allowed.
Overlapping Entries	Check this box to alert the system to allow multiple instances of key pairings, which keeps the previous session with the same key field open. For example, you might check this box if the list is tracking activity for an asset that supports multiple user logins.
In Memory Capacity (x1000)	This setting indicates the maximum number of session entries the system keeps in memory. The default value is 10,000. For most cases, 10,000 is appropriate; however, you may wish to adjust this setting if the devices you are monitoring for this session list contain a lot of data to ensure you have adequate memory cache available. As a best practice, be sure to set In Memory Capacity higher than the number of live sessions you anticipate. This helps optimize performance and, therefore, keeps results reliable.
Entry Expiration Time	Enter an expiration time in hours, minutes, and seconds for session list entries. This indicates the time after which entries are marked as terminated (if no explicit termination event is received previous to this). Maximum expiration is 24 days. The default is Unlimited, which means the entry never expires. An entry with no expiry date/time can only be terminated explicitly through user action on ArcSight Console, rule actions, or archives.
TTL Days	Set the <i>least</i> number of days a closed session should remain on the list before it is removed. Default is 0 days. Use 0 to keep the closed session indefinitely. The maximum number of days is 999999 .
Case Sensitivity	You can optionally configure the list to be case-sensitive or -insensitive. Furthermore for case-insensitive lists, you can specify case-insensitivity for keys only, or for both keys and values. The feature enables you to store and look up values in lists regardless of case. Select one: • Case-Sensitive (the default) • Key Case-Insensitive • Key & Value Case-Insensitive Important: After you save the list, you cannot change this setting. If you want to revert the case sensitivity setting, define a new list instead. Caution: Lookups on case-insensitive lists will slow down query and active channel performance. Make sure your queries and variables (used by channels) get values from <i>case-sensitive</i> lists.
Common and Assign fields	Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see " Common Resource Attribute Fields " on page 725.

- Under the Name column, replace *<Enter Name>* with a descriptive name for each session parameter you want to track.

The name you enter here appears as a label in the session list and in the Variable pick list. Names can contain spaces, such as **User Name**. For a list of restricted characters in field names, see ["Field Naming Restrictions" on page 506](#).

Columns for Start Time, End Time, and Creation Time are pre-defined.

4. Enter the corresponding data type, sub-type, and mark as key field as required. Refer to the following table for guidance:

Session List Column Types and Subtypes

Type	Subtype
IP Address	This field supports IPv4 or IPv6 address. If the value is an IPv6 address, the resulting address is displayed in simplified format if applicable. For example, 2001:db8:0000:0000:0000 is displayed as 2001:db8::
Date	This Date field is used as a default Timestamp value for interval-type queries on session lists.
Double, Integer, or Long	Select the applicable numeric type. Note: Leave the Subtype column blank even if you see the selections. The numeric subtypes MIN, MAX, and SUM are not supported in session lists.
MAC Address	MAC address of the format consisting of six groups of two hexadecimal digits per group. Use hyphen (-) as separators. For example 01-00-5E-90-10-FF
Resource Reference	Any ArcSight Resource such as asset, report, actor, and so on.
String	This is optional for lists in general but required, along with a Date field, if your list is time partitioned.
Key field	Select one or more fields that must be unique to indicate a session start. In most cases, you would select at least two fields to make a key-value pair. For example, in the case of a DHCP login event, when a new IP and zone combination are written to the list, this indicates that a new session has started.

Database columns are defined after the session list is created. Column definitions cannot be added, removed, or changed once the new session list is saved.

5. Click **Apply**.

The Filter tab for the list becomes enabled.

6. Click the **Filter** tab in the Session List Editor and define a filter that limits the number of events to consider for the new session list.

Session lists without filters must evaluate every event, which can negatively affect performance. The Filter tab presents the Field Set selection panel. Session list filters are different from regular filter resources; they use different fields.

Session lists are often concerned with logins to specific machines. In this case, you would write a filter that would limit evaluation to IP address ranges of interest. By filtering out all events except those targeting IP addresses in the DHCP server's subnet, for example, you are effectively limiting session list evaluation to inside traffic, reducing the overhead of session list evaluation. Other uses of session lists suggest other installation-specific knowledge that can be used to create session list filters that restrict the number of events matched against the session list.



Note: Filters are used to improve session list performance by restricting the number of events that must be evaluated. Filters, such as DHCP IP address ranges, are installation-specific. Therefore, consider adding a filter to pre-defined session lists, such as /All Session Lists/ArcSight Foundation/Network Monitoring/DHCP, to improve performance.

7. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
8. Click **Apply** to save and continue editing or **OK** to save and close.



Tip: Use the **Add Entry** button in the Session List Editor to manually add entries to the current session list.

More information:

- ["Understanding Session Correlation" on page 565](#)
- ["Example: Using Session Lists to Correlate Session Data on User Logins" on page 569](#)

Editing Session List Entries

Procedures in this topic include adding and deleting session list entries and terminating an entry.

Entries are added or removed to a session list through rule actions (see ["Rule Actions Reference" on page 539](#)). You can manually add or delete entries as described in this topic.

Where: **Navigator > Resources > Lists > Session Lists tab**



Caution: A session list can contain only one entry with the same key and StartTime value (up to milliseconds). This is useful in preventing duplicate session entries from multiple rule firings and from scheduled rules that are also deployed in real-time.

To add a session list entry based on an existing entry:

1. Right-click a session list and select **Show Entries**.
2. In the session list grid view, right-click an entry that is similar to the entry you would like to add. Choose **Edit**. The Session List Entry editor appears in the Inspect/Edit window.
3. Click a row's **Value** column to make changes. The column type may limit the kind of data that can be entered.
4. Click **Add** to post the changed entry as a new one.

To add a new session list entry:

1. Right-click a session list and select **Edit Session List**. The Session List Entry editor appears in the Inspect/Edit window.
2. Click **Add Entry**.
3. Click a row's **Value** column to make changes. The column type may limit the kind of data that can be entered.
4. Click **Add** to save the new entry. The **Reset** button clears all values.

To delete a session list entry:

1. Right-click a session list and select **Show Entries**.
2. In the session list grid view, right-click the entry that you would like to delete.
3. Select **Delete**. Confirm the deletion by clicking **Delete**.

To terminate a session list entry:

1. Right-click a session list and select **Show Entries**.
2. In the session list grid view, right-click the entry you want to terminate and select **Terminate Session Entry**.
3. Enter the date and time for the session end time. Click the  button for a context menu containing relative times such as Now, 1 hour ago, 1 day from now, and so on.
4. Click **OK**.

Moving, Copying, or Deleting a Session List

Where: Navigator > Resources > Lists > Session Lists tab

To move or copy a session list:

1. Drag and drop a session list into another group.
2. Select **Move** to move the session list, **Copy** to make a separate copy of the session list, or **Link** to create a copy of the session list that is linked to the original session list.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete a session list

1. Right-click a session list and select **Delete Session List**.
2. Click **Delete** and confirm.

Exporting a Session List

Purpose: Export session list entries to a CSV file so you can manage the information external to the Console.

Where: Navigator > Resources > Lists > Session Lists tab

1. Right-click a session list and select **Show Entries**.
The data in the session list is displayed in the Viewer panel as session list details.
2. In the Viewer panel, select one or more entries.
3. Right-click and choose either **Export CSV - Visible Columns** or **Export CSV - All Columns**.

This opens a file browser.

4. Browse to the location where you want to save the exported data, enter a file name in the File Name field, and click **Save**.

The entries you selected for export are saved as a CSV file in the chosen location.

Field Naming Restrictions

The following information on field naming restrictions applies to both session lists and field-based active lists. When you enter a name for your field, you are creating a database column. Field names, therefore, are subject to character restrictions consistent with database column names. The following table lists the characters you must not use in field names.

Disallowed Characters for Field Names

Disallowed Character	Description
&	Ampersand
*	Asterisk
@	At
^	Caret or circumflex
:	Colon
.	Dot or period
=	Equals
!	Exclamation point
>	Greater than
-	Hyphen or dash
<	Less than
(	Parenthesis, left
)	Parenthesis, right
+	Plus
'	Single quote
/	Slash, forward
[	Square bracket, left
]	Square bracket, right
	Vertical bar

Chapter 17: Rules Authoring

This section explains how to use rules to correlate events in your environment.

Topics include:

- ["Designing Rules" below](#)
- ["Rule Types" on the next page](#)
- ["Creating or Editing Rules" on page 509](#)
- ["Enabling and Disabling Rules" on page 511](#)
- ["Specifying Rule Thresholds and Aggregation" on page 528](#)
- ["Managing Rule Actions" on page 532](#)
- ["Converting Rule Types" on page 550](#)
- ["Testing Rules" on page 551](#)
- ["Verifying Rules with Events" on page 553](#)
- ["Deploying Real-time Rules" on page 555](#)
- ["Scheduling Rules" on page 558](#)



Note: You can configure your ESM installation in compact mode; or in distributed mode where you can have a cluster of multiple correlators, aggregators, and other components. Even if you initially installed ESM in compact mode, you can convert your system to distributed mode. The same applies to those who upgraded from a previous ESM version. Previous ESM versions are essentially ESM in compact mode. This installation and configuration process is covered in either the [ESM Installation Guide](#) or the [ESM Upgrade Guide](#). See also [ESM 101](#) for conceptual information, and the [ESM Administrator's Guide](#) for distributed correlation management.

Throughout this section, you will see supplemental information specific to **distributed correlation** which will be flagged accordingly. Unless specified, most of the topics in rules processing apply to both compact and distributed mode.

See also ["Checking the Status of the Distributed Correlation Cluster" on page 77](#) for related information.

Designing Rules

Creating rules involves defining the events the rule evaluates, thresholds, and actions you want the rule to trigger. Conditions define which events trigger the rule, thresholds determine when a condition is met and a correlation event is generated, and actions state what responses are taken when a correlation event is generated.

To define rule events and conditions, thresholds, and actions, begin by determining:

- Which event occurrences do I want to be aware of? This determines what **events** this rule needs to monitor and the **conditions** to be tested.
- How many times do I want the event or events to occur and within what time frame? This determines the rule's **threshold**.
- What actions should automatically occur when an event is generated? When should those actions occur? This determines the rule's **actions**.

Before you create rules, determine which events you want to monitor. Be specific and as clear as possible. For example, monitoring all events from a Cisco Router would not be as useful as monitoring all denied events from that Cisco Router. In addition, the more conditions you add to a rule, the more specific the rule becomes. Use the ArcSight data fields to guide you in selecting and specifying conditions. For more information, see ["Data Fields" on page 910](#).

Rule Types

ESM provides the following rule types:

Type	Description
Standard rules	Include all features for rule creation such as one or more event aliases (has joins), field aggregation options, and rule actions based on different triggers. You can convert a standard rule to a lightweight or pre-persistent rule ("Converting Rule Types" on page 550).
Lightweight rules	Include a small set of features for rule creation for faster and simpler rule processing. A lightweight rule: • Has only one event alias (no joins). • Does not aggregate data fields, therefore, the Aggregation tab is disabled. • Executes a specific action only on the On Every Event trigger. • Only allows active and session list actions. See also Rule Actions Reference for additional details on the Active List and Session List rule actions. • Does not generate correlation or audit events, although failures are logged. • Is processed earlier in the flow than standard rules. • Can be converted to other rule types ("Converting Rule Types" on page 550).

Type	Description
Pre-persistence rules	Include a small set of features to enable basic event analysis and the setting of various event fields, therefore enriching these base events, before the events themselves are persisted in the database. A typical usage for this rule type would be for threat-level formula calculations. A pre-persistence rule: • Has only one event alias (no joins). • Does not aggregate data fields, therefore, the Aggregation tab is disabled. • Executes a specific action only on the On Every Event trigger. • Can only perform the Set Event Field action. The action is applied to incoming base events. Values of the modified fields are available to standard and lightweight real-time rules, which run during the post-persistence processing flow. • Does not generate correlation or audit events, although failures are logged. • Is processed earlier in the flow than lightweight and standard rules. • Cannot be scheduled or replayed, since events occurring in the past have already been persisted and can no longer be modified. • Can be converted to other rule types (see "Converting Rule Types" on page 550).

Managing Rules

Like other resources, the rule-management tasks include creating, changing, deleting, and deploying them.

Creating or Editing Rules

Before creating rules, determine which events you want to monitor. Be as specific and as clear as possible. For example, monitoring all events from a Cisco Router would not be as useful as monitoring all **denied** events from that Cisco Router. In addition, the more conditions you add to a rule, the more specific the rule becomes.

Use the ArcSight data fields to guide you in selecting and specifying conditions.



Caution: If you are editing a standard rule because you want to change its rule type, follow the instructions in [Converting Rule Types](#).

Where: Navigator > Resources > Rules

To create or edit a rule:

1. If you are creating a rule, right-click a group and select **New Rule | <Rule Type>**. See ["Rule Types" on page 508](#) for guidelines on rule types.

If you are editing a rule, right-click the rule and select **Edit Rule**.

2. On the **Attributes** tab, enter or change the name in the **Name** text field.

The name is restricted to 25 characters. Be as descriptive as possible. The name is stored in the Event Name data field and appears in the Event Name column on the grid view.

3. Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

4. Required: Define conditions on the **Conditions** tab following instructions in ["Specifying Rule Conditions" on page 513](#). You cannot save the rule without specifying conditions. Non-standard rules have restrictions (see ["Rule Types" on page 508](#) for details).

To view the full conditions for the **MatchesFilter** operator, click the **Summary** tab and then click the **Expand Filter** button to display the filter conditions for debugging.

Note that in this case, the display of the **MatchesFilter** full logic does not display the sub-filter of the matched filter. Full logic is displayed only for the first level of matched filter conditions.

5. For standard rules, add correlating events, specify thresholds and time windows to qualify events, and aggregate incoming event data based on matching fields on the **Aggregation** tab. See ["Specifying Rule Thresholds and Aggregation" on page 528](#).



Note: The Aggregation tab is enabled for standard rules only.

6. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
7. Click **OK**.

Moving or Copying Rules



Note: You cannot move or copy a pre-persistence rule into a rule group that has been scheduled.

Where: Navigator > Resources > Rules

1. Select the rule and drag and drop it into another group. When you copy a rule, the default state of that rule is disabled.
2. Select **Move** to move the rule, **Copy** to make a separate copy of the rule, or **Link** to create a copy of the rule that is linked to the original rule.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

Enabling and Disabling Rules

Purpose:

- Enable rules that you have verified as working properly.
- Disable rules if they are being triggered too many times so that system performance is affected. You may then troubleshoot if necessary, before re-enabling the rules again.

When a rule is enabled or disabled, information on this rule appears in the notes on the rule. This information includes details of the change, such as the user that changed the rule and when the change was made.



Tip: ESM profiles rule performance by measuring their evaluation time on a sampling basis. You can view these results from the Rules Status dashboard and from there, you can manually disable rules which you deem expensive.



Note: Only rules deployed in Real-time Rules show up in a live channel when they are triggered. Therefore, after you have created and verified rules and you are ready to deploy them on real-time events, move or copy the rules to your user folder under Real-time Rules as described in ["Deploying Real-time Rules" on page 555](#).

Where: Navigator > Resources > Rules

To enable rules:

Right-click the rule and select **Enable Rule**. The rule is displayed as enabled or on (🔥) in the Navigator.

To disable rules:

Right-click a rule and select **Disable Rule**. The rule is displayed as disabled or off (🔌) in the Navigator.



Note: If ESM has automatically disabled a rule, you must still manually disable the rule as described; otherwise the rule will continue to fire and will be automatically disabled in a circular manner. This process is resource intensive. See ["Automatically Disabled Rules " on page 1035](#) for details.

To disable rule components:

You can disable certain components of a rule, such as particular rule triggers or a rule actions associated with particular triggers. For information on this, see:

- ["Activating or De-activating a Rule Trigger " on page 534](#)
- ["Enabling or Disabling a Rule Action" on page 535](#)

Viewing Rules and Their Correlation Events

ESM generates correlation events when a rule action is executed. Only standard rules generate correlation events.

Purpose: To display a standard rule's correlation events in a channel.

Procedure:

1. On the navigator's **Resources** panel, expand **Rules**.
The top node displays the rules you yourself created.
2. Locate your standard rule of interest. If necessary, expand a rule group to display a list of rules.
3. Right-click that rule and select **Create Channel with filter**.
A temporary channel is created on the Viewer panel, displaying the correlation events that resulted from executing that rule's actions. If the rule has not triggered for any reason (for example, because no matching events are found), then the channel is empty.

For more information about correlation events, see ["Rule Actions" on page 870](#).

You can save this temporary channel if you want. After saving, you can display this channel on the Console and the ArcSight Command Center.

Deleting Rules

Where: Navigator > Resources > Rules

1. Right-click a rule and select **Delete Rule**.
2. Click **Yes** to confirm.

Specifying Rule Conditions

Purpose:

To find events that match the rule condition statements; and if matching events are found:

1. Trigger the rule action(s).
2. Generate a correlation event.

Definition of terms:

- **Base events** are the events that match the rule's conditions. They are also called correlated events.
- **Correlation event** is a system-generated audit event that caused the rule to trigger.

Creating Rule Conditions

Purpose: To continue the process of creating or editing a rule.

The rule's Conditions tab provides a default event alias, **event1**, which you edit and to which you add condition statements for evaluation.



Note: Standard rules can have multiple event conditions. Lightweight and pre-persistence rules are limited to only one condition.

Where: Navigator > Resources > Rules

To specify rule conditions:

1. In the Rules Editor, select the **Conditions** tab.
2. To edit the event alias (change its default name), right-click **event1** and select **Edit**. Enter a new name for the event alias in the text field and click **OK**.

Because rules can monitor numerous events, aliases should be unique and descriptive. For example, if monitoring Cisco Router denied events, **Cisco Router denied** could be the alias name. The name appears as a branch under the **Event conditions** tree.

3. To add a condition statement to the event alias using the Common Conditions Editor table (usage rules and features of this editor are described in ["Common Conditions Editor \(CCE\)" on page 887](#)):
 - a. Locate the event name you want to use in the condition statement.
 - b. Select the logical operator (for example, =) to be used for comparing values. If you need help, see ["Logical Operators" on page 1006](#) for descriptions.
 - c. Select the value from the drop-down list under the Condition column to use as the basis for comparison.



Note: If you want to use a global variable to set a value for the condition statement, click the **+/- Global Variables** button and then choose the global variable from the resource selector popup. The selected global variable will be added to the Common Conditions Editor table at the bottom of the Edit panel. See ["Global Variables" on page 1002](#) for more information.

4. To add resource-specific condition statements, see:
 - ["Adding Filter Conditions" on the next page](#)
 - ["Adding Asset Conditions" on page 516](#)
 - ["Adding Vulnerability Conditions" on page 516](#)
 - ["Adding Active List \(InActiveList\) Conditions" on page 517](#)
5. For standard rules only: To add more event aliases, select **Event conditions** and click the **New Event Definition** button; or right-click **Event conditions** and select **New Event Definition**. Enter an event name in the **Alias** text field and click **OK**.

If you have more than one event alias, a **Matching Event** branch appears. This enables you to define a join relationship on the multiple event aliases. For more information on joining two events, see ["Creating Matching or Join Conditions" on page 519](#).

If you are working on a non-standard rule, you will not be able to save the rule if you have more than one event condition.

6. On the Conditions tab, click **Apply**.

The rule with the default threshold and action is created and listed in the Rules resource tree.



Note: The rule conditions are shown on the rule's Notes tab for historical purposes. For imported rules or rules created in previous versions, the Notes tab is updated only when the conditions are edited after the import or after the upgrade.

For standard rules only, see ["Specifying Rule Thresholds and Aggregation" on page 528](#) for aggregation time-frame options.

Adding Filter Conditions

Purpose: Find an event that matches, or does not match, an event in the specified filter. If found, generate a correlation event.

For more information on filters, see ["Filtering Events" on page 306](#).

To add a filter condition to a rule:

1. In the Rules Editor, select the **Conditions** tab and select the event alias to which you want to add a filter condition.
2. Click the **And**, **Or**, or **Not** button; or right-click a logical operator and choose **New Logical Operator**, then **And**, **Or**, or **Not**.

If there are existing conditions, you can tie them to the filter condition with either the AND, OR, or NOT logic operator. If AND is used, all the existing conditions and the filter condition must occur in the event. If OR is used, either the other existing conditions or the filter condition must occur. If NOT is used, all but the filter condition must occur.

3. Right-click the logical operator and select **New Matches Filter Condition**.
4. In the Filter Selector, select a filter and click **OK**.
5. On the Conditions tab, click **OK**.

The Common Condition Editor's buttons and commands are discussed further in ["Creating or Editing a Filter" on page 306](#).

See also ["Logical Operators" on page 1006](#), ["Condition Tree Command Buttons" on page 890](#), ["Condition Tree Context Menu Commands" on page 892](#), ["Common Conditions Editor \(CCE\)" on page 887](#), and ["Adding Conditions" on page 895](#).

Adding Asset Conditions

Purpose: Find an event, and if the specified asset is the source or target, generate a correlation event.

Assets are part of your network model as described in ["Modeling the Network" on page 119](#).

To add an asset condition to a rule:

1. In the Rules resource tree, right-click a rule and choose **Edit Rule**.
2. In the Rules Editor, select the **Conditions** tab.
3. Click the **And**, **Or**, or **Not** button, or right-click a logical operator and choose **New Logical Operator**, then **And**, **Or**, or **Not**.

If there are existing conditions, you can tie them to the asset condition with either the AND, OR, or NOT logic operator. If AND is used, all the existing conditions and the asset condition must occur in the event. If OR is used, either the existing conditions or the asset condition must occur. If NOT is used, all but the asset condition must occur.

4. Select the logical operator and click the **Assets** button on the rule editor toolbar, or right-click the logical operator and select **New Assets Condition**.
5. In the Assets panel below, select **Source Asset ID** to monitor if an asset is the source of an event or **Target Asset ID** to monitor if an asset is the target.
6. Select an asset or group and click **Apply**.

The asset condition appears in the Correlate section and is tied to any existing condition statements with the logic operator selected.

7. On the Conditions tab, click **OK**.

See also ["Logical Operators" on page 1006](#), ["Condition Tree Command Buttons" on page 890](#), ["Condition Tree Context Menu Commands" on page 892](#), ["Common Conditions Editor \(CCE\)" on page 887](#), and ["Adding Conditions" on page 895](#).

Adding Vulnerability Conditions

Purpose: Find an event that has the specified vulnerability, and if found, generate a correlation event.

For more information on vulnerabilities, see ["Modeling the Network" on page 119](#).

To add a vulnerability condition to a rule:

1. In the Rules resource tree, right-click a rule and choose **Edit Rule**.
2. In the Rules Editor, select the **Conditions** tab.
3. Click the **And**, **Or**, or **Not** button or right-click a logical operator and choose **New Logical Operator**, then **And**, **Or**, or **Not**.

If there are existing conditions, you can tie them to the vulnerability condition with either the AND, OR, or NOT logic operator. If AND is used, all the existing conditions and the vulnerability condition must occur in the event. If OR is used, either the existing conditions or the vulnerability condition must occur. If NOT is used, all but the vulnerability condition must occur.


4. Choose the logical operator and click the **Has Vulnerability** button on the rule editor toolbar, or right-click the logical operator and choose **New Has Vulnerability**.
5. In the **Vulnerability Selector**, select a vulnerability and click **OK**.

The vulnerability appears on the Conditions tab and is tied to any existing condition statements with the logic operator selected.

6. On the Conditions tab, click **OK**.

See also ["Logical Operators" on page 1006](#), ["Condition Tree Command Buttons" on page 890](#), ["Condition Tree Context Menu Commands" on page 892](#), ["Common Conditions Editor \(CCE\)" on page 887](#), and ["Adding Conditions" on page 895](#).


Adding Active List (InActiveList) Conditions

Use the Active List selector  to specify an active list that contains the argument for a condition. This condition evaluates whether an item or list of items is in an active list. You can use this to map a field or a global variable in the event schema to a corresponding field in an active list. It does not evaluate items in other non-event schemas (such as cases or assets).

Comparing values in two lists:

When the InActiveList condition is used to compare values in two lists, an additional option is shown where you can specify whether **All values in list field must match**.

For example, suppose you have a fields-based multi-mapped active list that has User Name as a key field and accepts entries with multiple roles for the same user in the Role Name field:




The 'Inspect/Edit' window for 'Active List: User Roles' shows a list of attributes and their values. The 'Active List' section includes fields like Name, Optimize Data, Capacity, TTL Days, TTL Hours, TTL Minutes, Allow multi-mappings, Partially cached, Time partitioned, and Case-Sensitivity. The 'Common' section includes Resource ID, External ID, Alias (Display Name), and Description. Below the list, there are radio buttons for 'Data' (Event-based or Fields-based) and a checkbox for 'Key Fields'. A table at the bottom lists fields with their types and sub-types.

Name	Type	Sub-type	Key-field
User Name	String		☒
Role Name	String		☐

Then suppose you set up an Active List (InActiveList) rule condition to compare the value of Role Name to a list type string field, like ActorByAccountID.FullName. If you then get list entries in your active list (for example, user “Samantha Stevens” with roles as both “Administrator” and “Development Lead”), then your rule results in a comparison of two lists:

- The list of Samantha Steven’s roles
- The ActorByAccountID.FullName list



The 'Rule Editor: New Rule' window shows the 'Conditions' tab. Under 'Event conditions', there is a condition named 'event1' with a sub-condition 'InActiveList(/All Active Lists/Personal/admin's Active Lists/User Roles)'. Below this, the 'InActiveList' configuration is shown, including a dropdown for 'User Roles' and a table for field comparisons.

Name	Field
* User Name	ActorByAccountID.Full Name
* Role Name	ActorByAccountID.Role Name

☒ All values in list fields must match

Buttons: Test, OK, Cancel, Apply, Help

- If **All values in list field must match** is checked (selected), the Active List condition evaluates to true only if all values in both lists match (that is, all values must be in both lists for the condition to be true).
- If **All values in list field must match** is *not checked* (de-selected), then if *any* field matches (is in both lists), the condition statement evaluates to true. This is the default behavior for queries.



Note: When using the InActiveList condition, remember the following:

- The InActiveList condition evaluates single-value and multi-value attributes. The field you map could possibly return multiple values (for example, a user could have multiple roles). In the case of multi-value attributes, if any one value matches, the condition evaluates to true.
- A condition that tests for whether all or any values in a list match is only available to specify on in-memory operations (for example, in rules, filters, data monitors).

See also ["Logical Operators" on page 1006](#), ["Condition Tree Command Buttons" on page 890](#), ["Condition Tree Context Menu Commands" on page 892](#), ["Common Conditions Editor \(CCE\)" on page 887](#), and ["Adding Conditions" on page 895](#).

Creating Matching or Join Conditions

This topic applies to standard rules only. It provides examples for creating matching or join conditions and suggestions for optimizing the use of resources to process such rules.

About:

A matching or join condition is a condition statement that joins two data fields with the Matching or Join condition logic operator on the Conditions tab. Creating matching or join conditions using data fields provides the flexibility of creating conditions without knowing the specific data field's values. You can create the following join data field conditions:

- **Same data field for two events** use this format: EventOne <data field A> <logic operator> EventTwo <data field A>. For example, EventOne Source Address = EventTwo Source Address. In this example, both event data fields must have the same value. This rule is useful when monitoring activity from an unknown Source Address that is generating numerous events.
- **Different data fields for two events** use this format: EventOne <data field A> <logic operator> EventTwo <data field B>. For example, EventOne Source Address = EventTwo Target Address. In this example, the Source Address of the first event must equal the Target Address of the second event.
- **Different data fields for the same event** use this format: EventOne <data field A> <logic operator> EventOne <data field B>. For example, EventOne Source Address

= EventOne Target Address. In this example, the Source Address must equal the Target Address of the same event.



Note: There is a relatively high memory cost for join rules with low-selectivity join conditions (such as same source IP address or same target IP address). Just like SQL queries, the more selective the conditions (the conditions on the individual events as well as the join conditions), the less expensive it is to execute, because fewer conditions match.

You can reduce the correlation engine's memory consumption by as much as 50% in some cases through some techniques. When authoring a rule, you order conditions on the events to be correlated (or joined) by placing the most restrictive conditions first; for example, adding join conditions like event1's Source Address = event2's Source Address or event2's Detect Time = event1's Detect Time.

If your condition specifies more than one event alias, you can set any or all of them with the **Consume After Match** flag. This means that if a matching event is found and the rule is triggered, the rule will not correlate the event any further. Without the Consume After Match flag, the event is kept in working memory even after a matching event is found and the rule has been triggered. The event alias continues to be combined with events matching other aliases until the event itself expires.

If enabled, the Consume flag appears next to the event alias on the Conditions tab:

EventOne (Consume after match)



Tip: See also ["Optimizing the Evaluation of Event Conditions" on page 524](#).



Note: Lightweight and pre-persistence rules have only one event, therefore, the Consume After Match option is not available.

To create a rule with matches or joins (with two or more events):

1. In the Rules resource tree, right-click a rule and choose **Edit Rule**.
2. In the Rules Editor, select the **Conditions** tab.
3. Select the **Matching Event** branch and:
 - a. Select **New Logical Operator**.
 - b. Select **And, Or, or Not**.
 - c. Add the second event that is tied to the first event.

When adding join conditions, you need to decide how the new condition ties to the existing events in the rule. If you use **And**, the new join condition must occur, in addition to

the existing events, to trigger the rule. If you use **Or**, the new join condition or the existing events must occur. If you use **Not**, all but the new join condition must occur. The logical operator appears as a branch under Joins.

4. Click the  (**Join Condition**) button or right-click the logical operator and select **New Join Condition**.

A condition statement appears, displaying event, data field, and logic operator text fields. These fields are combined to create *<event> <data field> <logic operator> <event> <data field>* condition statements. For example, if monitoring for the same Source Address data field in EventOne and EventTwo, the condition statement would be EventOne Source Address = EventTwo Source Address.

5. Select one of the following join data field conditions to use in the following steps:
 - When monitoring for the same data fields for two events use EventOne *<data field A><logic operator>* EventTwo *<data field A>*.
 - When monitoring for different data fields for two events use EventOne *<data field A><logic operator>* EventTwo *<data field B>*.

6. In the text fields, select an event and data field from the drop-down menus.

Select data fields that you want to monitor but for which you do not have values. For more information, see ["Data Fields" on page 910](#).

7. Select a logic operator from the drop-down menu.
8. Select an event and data field from the drop-down menus.
9. Optionally right-click and select **Consume After Match** on one, some, or all of the event aliases.

Doing so reduces the number of rule firings by using the matching event in only one join.

10. Click **OK**.

The join data field condition appears as a branch under the Matching Event logical operator.

11. On the Conditions tab, click **OK**.

See also ["Logical Operators" on page 1006](#), ["Condition Tree Command Buttons" on page 890](#), ["Condition Tree Context Menu Commands" on page 892](#), ["Common Conditions Editor \(CCE\)" on page 887](#), and ["Adding Conditions" on page 895](#).

Editing or Deleting Join Data Field Conditions

For related information, see ["Creating Matching or Join Conditions" on page 519](#).

Where: Navigator > Resources > Rules

1. Right-click a rule and select **Edit Rule**.
2. In the Rules Editor, select the **Conditions** tab and do the following:
 - To edit the logical operator, right-click the logical operator and select **Edit** or select the logical operator and press **Enter**. In the text field, select a logical operator and click **OK**.
 - To edit the condition statement, right-click the condition statement and select **Edit**, or select the condition statement and press **Enter**. In the text field, make your edits and click **OK**. For more information, see ["Creating or Editing Rules" on page 509](#).
 - To delete the Matching Event event, right-click **Matching Event** and select **Delete**. In the popup, click **Yes**. The event, its logical operators, and condition statements are deleted.
 - To delete the logical operator, right-click the logical operator and select **Delete**. In the popup, click **Yes**. The logical operator and all its condition statements are deleted.
 - To delete the condition statement, right-click the condition statement and select **Delete**. In the popup, click **Yes**.
3. Click **OK**.

Negating Event Conditions

This topic applies to standard rules.

Purpose: To catch events that you expect to happen in a sequence of events, but events do not happen after all. You do this by negating the expected event in the rule's Conditions tab.

Negated events depend upon other events that have happened. For purposes of discussion, let us refer to these events that happen as positive events.

Where: Navigator > Resources > Rules

Prerequisite:

The rule must have two or more event conditions, so that you can negate at least one. To create event conditions, see ["Creating or Editing Rules" on page 509](#).

Scenario 1: Monitoring past events to catch a non-occurring event

This scenario shows two expected events that must appear in sequence.

1. Someone physically accesses a system (in the rule, call it BadgeScan event).
2. Someone accesses an application (in the rule, call it Login event).

Your rule conditions specify that BadgeScan must occur first before Login. You want the rule to trigger if these events are not received in that sequence. In this case, you negate the BadgeScan event. Both events must have occurred (they are past events) before the rule triggers.

Scenario 2: Monitoring a future non-occurring event

In this scenario, you negate a *future* event condition. For example, consider this sequence of events you want to monitor:

1. A server reboots (ServerReboot event).
2. The server successfully comes up and is available again (ServerUp event).
3. If the server does not come up, you want to be notified.

In this case, you will negate the ServerUp event condition so that the rule is triggered if that event is *not received* (the server does not come up from a reboot) on the same device.

A *time out* property is used in conjunction with negating an event condition. If the negated event is not received within the specified timeout, then the rule is triggered. For example, you can configure your rule to notify you if a server that reboots does not start up successfully.

To negate event conditions:

1. Right-click a standard rule and select **Edit Rule**.
2. In the Rules Editor, select the **Conditions** tab.
3. Right-click the event alias of interest and select **Negated**.
4. Right-click the negated event alias and select **Set Negated Alias Timeout**.
5. In the popup, enter a time out value in seconds, minutes, or hours.

Time Out is the amount of time to wait between the occurrence of the positive event and the non-occurrence of the negated event, after which the rule is triggered. This value is required.

Use these **important guidelines** for setting the time out value:

- The [Alias Expiration](#) time for all positive event conditions must be greater than the negated event condition's time out value.
- Time out values are cumulative. The rule will wait for the sum of all event timeouts before firing, therefore, the aggregated time must be greater than the negated event

condition's time out value. See the description of the Time out setting in ["Aggregation Time Criteria" on page 530](#).

- If a rule has multiple negated event aliases, set the timeout of one negated alias to **three minutes**, then set the remaining timeout values to zero. For example, consider a rule with three event aliases: event1 is positive, event2 is negated with timeout = 1 minute, and event3 is negated with timeout = 2 minutes. The rule will not trigger until at least 3 minutes after event1 has been matched. Moreover, if the event expiration time (by default the aggregation time window) is only 2 minutes, the rule will not trigger at all because event1 will be removed from memory prior to the cumulative timeout.

On the Conditions tab, the negated event is preceded by an exclamation point (!) and the time out period appears next to the event. The following example shows a five-minute time out period.

```
!<EventAliasName> (Time Out: 5m)
```


6. Click **OK** to save the time out value.
7. To remove the Negated flag, right-click the negated event and select **Negated** again.

See also ["Logical Operators" on page 1006](#), ["Condition Tree Command Buttons" on page 890](#), ["Condition Tree Context Menu Commands" on page 892](#), ["Common Conditions Editor \(CCE\)" on page 887](#), and ["Adding Conditions" on page 895](#).

Optimizing the Evaluation of Event Conditions

This topic is written for advanced content authors. The topic describes how to automate the optimization of event conditions to reduce impact on CPU usage.

Evaluating event conditions is one of the most resource-intensive operation in event processing. ESM evaluates event conditions in the sequence they appear on the rule's Conditions tab. The following example shows event conditions:



The outer and inner nodes indicate nested conditions. As a guideline, order these conditions from the most economical to the most expensive by putting the least costly condition on top. However, this may not be enough; you must also consider the TRUE/FALSE rate and whether you are using OR or AND operators.

Related topics:

- ["Automating Condition Optimization" below](#)
- ["Disabling the Optimization Feature" on the next page](#)
- ["Tracing the Optimization" on the next page](#)

Automating Condition Optimization

Requirement: To optimize the evaluation of event conditions, rules must be deployed in the Console's Real-time Rules.

About:

ESM has the ability to evaluate event conditions and if necessary, re-sequence the conditions to be evaluated through a property setting.

```
rule.dm.optimize.evaluation=true
```

- In compact ESM, the setting is defined in `server.properties` file.
- In distributed mode ESM, in addition to `server.properties`, the setting is defined in the correlator node's `correlator.properties` file. Depending on your number of correlator instances, you should replicate the same setting in each node's `correlator.properties`.

The reason why the same property needs to be in both properties files is that standard and lightweight rules are affected by `correlator.properties`; while pre-persistence rules are affected by `server.properties`.

This setting automates the process of evaluating and optimizing all of your deployed rules. The optimization process occurs in memory and will leave your resources intact. If there are changes, for example, real-time rules are updated, or the user de-activates and then re-activates the rule, optimization runs again.

One or more log files record the activities of the optimization process. For compact ESM, the file is `server.log`. For distributed ESM, the log files are `server.log` and `correlator.log` files.

The following line in the applicable log file denotes the end of optimization:

```
End of conditions optimization in <n> events
```

See the "Configuration" section in the [ESM Administrator's Guide](#) for more information:

- For detailed instructions on how to add settings to properties files, refer to the topic, “Managing and Changing Properties File Settings.”
- For information on the `server.log` file, refer to the topic, “Configuring Manager Logging.”

Disabling the Optimization Feature

Optimization runs only once on all existing rules, then the profiles for the evaluated rules are stored in the log file. However, if you changed any deployed rule, the changed rule is evaluated and optimized again. Changes include adding and updating a rule, and disabling and re-enabling a rule.

When you no longer need the optimization feature, change the optimization setting to false in `server.properties` and, additionally in distributed correlation mode, `correlator.properties` (in all correlator instances if you have more than one).

Tracing the Optimization

The tracing feature described here is optional. However, if you want to use this feature, you must turn off the condition optimization setting. The tracing feature enables you to capture information about how your rules were optimized. The information is stored here:

- `server.log` if you have ESM in compact mode
- `server.log` and `correlator.log` if you have ESM in distributed mode

To use the tracing feature, use this property setting:

```
rule.dm.trace.optimize.evaluation=true
```

Enter the above setting in the following files:

- In compact ESM, edit the `server.properties` file.
- In distributed ESM, edit `server.properties` and `correlator.properties` files.

This setting records in the log files the original sequence of conditions and how these conditions are re-ordered (optimized).

To locate the information in the log file, search for the rule’s name. The following example shows the log statements that include the evaluated rule’s URI (shown to be under Real-Time Rules) and a sampling of the profiled values. The profile describes the average time in nanoseconds a specific condition took to process. The condition’s position is indicated by an index starting with position 0. The profile in the log indicates that the costliest condition at Index 4 took an average of 2,375 nanoseconds over a sampling of 23:

```
Standard Rule
[2014-06-25 10:15:56,630][INFO
][default.com.arcsight.common.simplerulesengine.engine.d] Optimizing /All Rules/Real-
time Rules/Intrusion Monitoring/Attack Monitoring/DoS/SYN Flood Detected by IDS or
Firewall
Profiled values are
[ Index=2, avgTime=366 ,numSamples=23, Index=3, avgTime=407 ,numSamples=23, Index=1,
avgTime=434 ,numSamples=23, Index=0, avgTime=928 ,numSamples=22, Index=5,
avgTime=1793 ,numSamples=23, Index=4, avgTime=2375 ,numSamples=23]
```

The following lines indicate the order of conditions (denoted as nodes) as originally defined. Red numbers correspond to the index:

```
Original nodes are
Node : name Contains "syn" 0
Node : type EQ "Base" 1
Node : categorySignificance StartsWith "/Compromise" 2
Node : categoryTechnique EQ "/DoS" 3
Node : Not (<deviceEventClassId, name, attackerZoneName, attackerAddress,4
targetZoneName, targetAddress>InActiveList( "H7nDOZB0BACB337GCqa7qw==:/All Active
Lists/ArcSight System/Tuning/Event-based Rule Exclusions::" ) )
Node : ( ( categoryDeviceGroup StartsWith "/IDS" And type EQ "Base" And 5
categoryDeviceGroup NE "/IDS/Host/Antivirus" ) Or categoryDeviceGroup EQ
"/Firewall" )
```

The following lines indicate the optimized order of conditions (denoted as nodes):

```
Optimized nodes are
Node : categorySignificance StartsWith "/Compromise" Moved from position 2
Node : categoryTechnique EQ "/DoS" Moved from position 3
Node : type EQ "Base" Moved from position 1
Node : name Contains "syn" Moved from position 0
Node : ( ( categoryDeviceGroup StartsWith "/IDS" And type EQ "Base" And
categoryDeviceGroup NE "/IDS/Host/Antivirus" ) Or categoryDeviceGroup EQ
"/Firewall" ) Moved from position 5
Node : Not (<deviceEventClassId, name, attackerZoneName, attackerAddress,
targetZoneName, targetAddress>InActiveList( "H7nDOZB0BACB337GCqa7qw==:/All Active
Lists/ArcSight System/Tuning/Event-based Rule Exclusions::" ) )Moved from position 4
```

Observe how nodes were moved up and down to indicate which conditions are evaluated first. Index 4, profiled as the costliest, is now evaluated last.

To save a rule in its optimized state:

If you are a content author, you can leverage this log information to modify your rule permanently. This procedure is optional.

1. Refer to the log's Optimized nodes section.
2. Edit the rule by re-ordering the rule conditions to match the log's optimized sequence.
3. Save the rule in its optimized state.

You may then disable the optimization tool, described in ["Disabling the Optimization Feature" on the previous page](#).

Specifying Rule Thresholds and Aggregation

Thresholds are defined as an aggregate number of occurrences within a time span. When a threshold is met, the rule triggers an action.

A rule's threshold is defined in the Rules Editor's Aggregation tab. This tab is enabled for standard rules only. It is disabled for lightweight and pre-persistence rules.



Caution: If you set a rule to aggregate over fields of a *multi-mapped active list* or *overlapping session list*, the rule might fire multiple times, once for each field value in the corresponding list entries. The Console displays a warning when such a list field is selected in the Aggregation tab.

Do not set rules to aggregate over multi-mapped active list or overlapping session list fields, *and also* add entries to the same list in a rule action ("[Adding, Editing, or Removing a Rule Action](#)" on [page 533](#)). Setting *both* aggregation and rule actions to add entries to the same multi-mapped or overlapped list can cause the number of triggered rules to increase to an unmanageable level.

Setting or Changing Rule Thresholds

Purpose: To continue the process of creating or editing a rule. Here, you will specify the number of events to be matched within a specified timeframe.

1. In the Rules Editor, select the **Aggregation** tab.
2. In the **Number of Matches** field, enter a number if you want the rule to match more than one event.
3. In the **Time Frame** field, enter a number and select a time unit. For example, enter **2** and select **Minutes**.
4. If you want to aggregate on the basis of certain fields' content being distinct, click **Add** under the **Aggregate only if these fields are unique** pane to select the fields to use. Select fields from global variables, field sets, and local variables.



Tip: Fields are *unique* only when the combined value of all fields is unique. For example, suppose you wanted to aggregate on three fields: Event Name, Event Message, and Category Outcome, with a threshold of two matches. If you received two events both with values of Failed Login, Attempt, and Failure for these fields, respectively, these events *would not* be aggregated.

However, if you received only one event like the first example, and another with values of Failed Login, Attempt, and **Success**, these two events *would be* aggregated because the combined value is not the same for the given threshold number of events.

Aggregating on unique fields is applicable when you want to monitor widespread conditions, such as an attack on ten unique systems.

You can use the rule action to set an event field with a unique aggregation field's value. See [Set Event Field](#) for details.

5. If you want to aggregate on the basis of certain fields' content being identical, click **Add** under the **Aggregate only if these fields are identical** pane to select the fields to use. Select fields from global variables, field sets, and local variables.
6. Click **OK**.

The choices you make are expressed as a conditional statement in the **Summary** panel.

Examples of Grouping Unique or Identical Field Values

You can use aggregation techniques to group unique or identical field values and map them into an active list through the Add to Active List rule action (refer to the rule action, [Set Event Field](#)).



Note: When unique field aggregation is used, all list actions (add, remove, terminate) are fired once for each unique value (or set of values). How these values impact the list depends on the configuration of the list and how the unique fields map into that list.

In the following examples, the list has SourceAddress and SourceUserName as key columns.

However, If the list has no key columns or multi-mapped active lists are used, then all columns are functionally keys.

In the **Case 1** example below, an AddToActiveList action would be performed twice:

- If the unique values map to keys in the list (SourceAddress and SourceUserName), the result would be 2 entries, each with a count of 1.
- If the list key is TargetAddress, the result would be 1 entry (key=2.2.2.2) with a count of 2.

For the following examples, assume there is an event-based active list that maps the following:

IP Address = Source Address

Name = Source User Name

Consider a set of events with the following values:

SourceAddress	SourceUserName	TargetAddress
1.2.3.4	sumerian	2.2.2.2
1.2.3.4	agta	2.2.2.2
1.2.3.4	sumerian	2.2.2.2
1.3.5.7	trojan	2.2.2.2
1.3.5.7	agta	2.2.2.2

Case 1: Unique aggregation on one field

You would like to capture the unique source addresses. The fields in your Aggregation tab would be something like:

Aggregate only if these fields are unique: SourceAddress

Aggregate only if these fields are identical: TargetAddress

After aggregation and through the Add to Active List rule action, the active list entries would consist of:

IP Address	Name
1.2.3.4	sumerian
1.3.5.7	agta

Case 2: Unique aggregation on two fields

Using the same event set, this time the fields in the Aggregation tab would be:

Aggregate only if these fields are unique: SourceAddress, SourceUserName

Aggregate only if these fields are identical: TargetAddress

With the Add to Active List rule action, the active list entries would consist of:

IP Address	Name
1.2.3.4	sumerian
1.2.3.4	agta
1.3.5.7	trojan
1.3.5.7	agta

Aggregation Time Criteria

The ArcSight Console provides time-evaluation criteria that can affect event-occurrence aggregation and rule triggering. You apply these to rules through the Aggregation tab and the

statement panel of the Conditions tab.

Aggregation is based on an event's End Time value, not Manager Receipt Time. However, events are not kept in memory indefinitely, therefore if some events are received after a long delay (such as an hour or so), they will not be matched with events that have already been removed from memory.



Note: Aggregation Time in Distributed Correlation

If you have join rules, rules with negated event aliases, or non-join rules based on thresholds, you might observe that the rules will not trigger. This is due to aggregation time that is too short. If the events arrive at the aggregator at different times, it is possible that the first matching event has expired by the time the second matching event arrives.

Workaround: If rules do not trigger, increase the aggregation window time (Time Frame option).

Related topics:

- ["Creating Matching or Join Conditions" on page 519](#)
- ["Negating Event Conditions" on page 522](#)
- ["Setting or Changing Rule Thresholds" on page 528](#)

Aggregation Time Criteria

Criteria	Application
Time Frame	Set on the Rule Editor's Aggregation tab, Time Frame establishes the time span for occurrence aggregation. Event-occurrence aggregation is always controlled by Time Frame. Secondly, Time Frame becomes the default for global and alias expiration time, if these are not set separately. Note: You can set the Rule Action trigger On Time Unit in conjunction with the Aggregation Time Frame to limit the number of times a rule is triggered. See "Threshold Triggering Options" on page 535.
Global Expiration	Set on the Conditions tab, a global expiration applies to an entire rule. This is the amount of time that qualifying events for all aliases are retained in memory for evaluation, based on Manager receipt-time. Setting an alias expiration overrides a global expiration, if present. To set Global Expiration, right-click the rule's root node (Correlate) in the Conditions tab and choose Set Global Expiration Time.

Aggregation Time Criteria	
Criteria	Application
Alias Expiration	Set on the Conditions tab, an alias expiration applies to a single event alias within a rule. This is the amount of time that, for this alias only, a qualifying event is retained in memory for evaluation, based on Manager receipt time. Setting an alias expiration overrides a global expiration, if present. To set Alias Expiration, right-click an event alias in the Conditions tab and choose Set Alias Expiration Time. An event with an expiration time is displayed with an indicator, for example: event1 (Wait time: 5m) To remove the alias expiration time, right-click the event alias and change the time to 0.
Matching Time	Set on the Conditions tab, a matching time creates a time-proximity comparison for multiple-alias rules, based on events' actual creation times. When two or more rule-condition aliases are present, a Matching Event node appears. You can right-click this node and choose Set Matching Time to require events' original timestamps (specifically, the event's original end-time) to fall within a range. Note that this time-proximity test is independent of and different than the memory-retention parameter set by global or alias expiration.
Time out	Set on the Conditions tab, you are prompted to set a time out value in seconds, minutes, or hours when you set an event alias to Negated. The time out begins after receipt of all positive events. If a negated event is not received within this time out period, then the rule is triggered. Note: If you have multiple negated events with different time out settings, the longest time out period takes precedence.

Deleting Aggregation from a Rule

1. In the Rules resource tree, right-click a rule and choose **Edit Rule**.
2. In the Rules Editor, select the **Aggregation** tab.
3. In the **Aggregate only if these fields are unique** or **Aggregate only if these fields are identical** lists, select the fields to delete and click **Remove**.
4. Click **OK**.

Managing Rule Actions


The Actions tab of the Rules Editor offers a consistent interface for defining actions to take based on the thresholds of the events that trigger them.

In the Actions tab, you click the buttons in the top row to Add, Edit, or Remove event-action sets for rules. Click **Hide Empty Triggers** to hide or show triggers not currently used.

See also ["Rule-triggering Timing" on page 1040](#) for more details.



Note: Rules, rule triggers, and rule actions can be enabled or disabled at various levels. The rule itself can be enabled or disabled, the trigger on a particular rule can be activated or deactivated, and a rule action associated with a particular trigger can be enabled or disabled. Details on rule triggers and rule actions are described in this topic. For more information and a summary, see also ["Enabling and Disabling Rules" on page 511](#).



Adding, Editing, or Removing a Rule Action

Where: Navigator > Resources > Rules:

To add or edit a rule action:

1. In the rule's editor, display the rule's **Actions** tab.

If you are adding a rule action for standard rules, the first trigger **On First Event** is active by default. For other rule types, only the **On Every Event** trigger is active and all other triggers are disabled.

For standard rules, select an applicable threshold trigger that is active. If the desired trigger is not active, right-click it and select **Activate Trigger**.

2. If you are adding an action, click **Add** (), then
 - For standard rules, choose an action from the available options.
 - For lightweight rules, choose either **Active List** or **Session List**.
 - For pre-persistence rules, only the **Set Event Field** action is available.

If you are editing an action, click **Edit** to open that action's Add Action dialog box.

3. If you are adding an action, in the **Add "Action Name" Action** dialog box, set the action's parameters, if present.

If you are editing an action, change the action's parameters as required.

See ["Rule Actions Reference" on page 539](#) for information about rule action types.

 **Tip:** You can use references to Velocity Templates as parameters for rule actions to derive values from event fields and variables. See ["Velocity Templates" on page 1088](#).


4. Click **OK** to add the new action to the rule's threshold trigger.

To remove a rule action:

Select an action below a trigger in the **Actions** tab and click **Remove**.

Activating or De-activating a Rule Trigger

The rule editor displays available triggers on the Actions tab:



See ["Threshold Triggering Options" on the next page](#) for descriptions.

For new rules, the first trigger is active by default. When a trigger is activated, all enabled rule actions it contains are triggered when conditions are met.

- To activate a rule trigger, select the trigger in the **Actions** tab and click **Activate Trigger**.
- To de-activate a rule trigger, select the trigger in the **Actions** tab and click **De-Activate Trigger**.

You can add rule actions to any trigger regardless of their state.

Enabling or Disabling a Rule Action

For finer-grained control over which rules are triggered when, you can enable or disable a rule action associated with any of the triggers.

- To disable an action, select an action below a trigger in the **Actions** tab and click **Disable**.
- To enable an action, select an action below a trigger in the **Actions** tab and click **Enable**.

Threshold Triggering Options

Consider the following factors for determining your triggering options:

- The minimum threshold value you can set is **1**.
- Triggering actions on every or subsequent occurrence can quickly use up resources. Use these options conservatively.
- For threshold-based triggers, only a single correlation event is triggered on receipt of any single incoming event, even if that event has an aggregated event count high enough to trigger multiple firings. This is by design to prevent excessive firings. For example, if a rule has a threshold of 10, an event with an aggregated event count of 200 triggers only one rule firing (not 20).


Trigger Thresholds

Trigger	Threshold
On First Event	The first time rule conditions are met, overriding aggregation threshold settings. This is the default trigger.
On Subsequent Events	The second and subsequent times rule conditions are met (not the first), overriding aggregation threshold settings.
On Every Event	Every time rule conditions are met, overriding aggregation threshold settings. Note: This is the only trigger available for lightweight and pre-persistence rules.
On First Threshold	For the number of matches greater than 1, the first time rule conditions and threshold settings are met.

Trigger Thresholds, continued

Trigger	Threshold
On Subsequent Thresholds	For the number of matches greater than 1, the second and subsequent times rule conditions and threshold setting are met, not the first.
On Every Threshold	Every time rule conditions and threshold settings are met.

Trigger Thresholds, continued

Trigger	Threshold
On Time Unit	Defines an action to take if the given threshold is met in the specified number of minutes specified. (When: On Time Unit: Every <NumberOfMinutes>). Notes: - With On Time Unit (OTU), the minimum threshold value you must set is 2. This setting can work in conjunction with aggregation to limit the number of times a rule is triggered. For example, aggregation is set to 2 matches in 1 minute and you get 50 matches in 1 minute (depending on how you set the rule actions). If you then specify the rule to trigger at On Time Unit = 1 minute, even if there were 50 matches in 1 minute, the rule would only trigger once per minute when the aggregation threshold is met. - The list of correlated events attached to the On Time Unit trigger excludes the events composing the first threshold. For example, if the threshold is 2 and 5 matching events are found, the first 2 events are excluded and only the remaining 3 are included in the list of correlated events. If you want to include the missing first two events for the threshold rule firing, you can additionally use these other triggers, On First Threshold or On Every Threshold in conjunction with On Time Unit. In this case, you will not see the first two events as part of On Time Unit. Instead, the first two events will be part of On First Threshold or On Every Threshold. - Activating On Time Unit does not imply that a rule is triggered on the first event, on subsequent events, or on every event that meets conditions. This specifically sets the rule to trigger for every given On Time Unit <i>if</i> aggregation thresholds are met. - Be sure to set On Time Unit to less than or the same value as the aggregation Time Frame to prevent getting an extra correlation event for the rule itself.
On Time Window Expiration	Expiration time of threshold settings When the On Time Window Expiration (OTWE) trigger is activated, it includes an option to display a cumulative rule chain (a summary of triggered rules) at the end of the triggered rules list. By default, the cumulative rule chain option on an activated OTWE trigger is off. To toggle the option between On and Off, right-click the <i>active</i> OTWE trigger and select On or Off on the cumulative rule chain option as needed.  When an OTW trigger activates a rule, a correlation event is generated. If the cumulative rule chain option is <i>on</i>, the correlation event contains all the base events from the first threshold to the time window expiration. If the cumulative rule chain option is <i>off</i>, the generated correlation event contains events from the last threshold to the time window expiration. Limitation: Unique aggregation does not work with the On Time Window Expiration trigger if cumulative rule chain set to <i>on</i>. See "Setting or Changing Rule Thresholds" on page 528 for information on unique aggregation in rules.

Rule Actions Best Practices

For rule actions, consider the following factors:

- **Action sequence**

Add actions in the order in which you want them to be executed. For example, to set a static value in an active list, first add the action, **Set Event Field**; then add the action, **Add to Active List**.

The Editor does not always match the internal representation of the specified order of rule actions. However, if you add rule actions in the proper order, that order is maintained internally.

Actions added to a rule show up the first time in the order you add them. You can continue to modify these and they show up in this order. After you click **Apply**, the display reorders the actions so that **Add to Active List** shows up first even though the internal representation has not been modified. Even so, rule actions continue to work as expected unless you change the order. For example, if you delete the Set Event Field action then add it back in after Add to Active List action is already configured, the rule actions are mis-ordered and do not trigger as expected.

- **Rule actions for lightweight and pre-persistence rules**

If you are creating or editing a lightweight rule, the rule can only act on active and session lists. If you are creating or editing a pre-persistence rule, the only available action is to set an event field.


- **Use of velocity expressions in rule actions involving lists**

- You can use references to Velocity Templates as parameters for rule actions to derive values from event fields and variables. (For additional details, see ["Velocity Templates" on page 1088](#).)

If you are using velocity expressions to derive values from variables and your rule is acting on an active or session list, perform these extra steps in conjunction with your action:

1. Aggregate over the fields of interest on the rule's **Aggregation** tab.
2. Use the **Set Event Field** action to set unused fields to the fields you specified for aggregation. Start with the **\$** symbol followed by the exact name of the variable but without any special characters like spaces and dots. For example, if the variable is **ActorByAccountID.Last Name**, you may use something like:

```
$ActorByAccountIDLast_Name
```



3. Continue by specifying the list to be acted on by the rule.



Note: Duplicate rule actions after a crash recovery:

If you stop ESM, it takes a checkpoint of the rules engine so that it knows what actions have been performed and where it stopped. If ESM crashes in such a way that it cannot take a checkpoint (during a power failure, for example), it returns to the last checkpoint when ESM restarts, and replays events from there. Any actions that occurred between that checkpoint and the ESM crash are therefore repeated. Repeated actions that generate audit events generate duplicate audit events.

You should investigate repeated actions that do not duplicate well. For example, if an action adds an item to an Active List, that item's counter will be incremented. If the action runs a command, it will run the command again, and so on.

Rule Actions Reference

The following table contains rule actions that are available if you right-click a trigger on a rule's Actions tab and select **Add**.

Rule Actions

Action	Expanded Menu Option	Description
Set Event Field		Fill in a data field value for correlation events generated by the rule using one of these methods: • Select from the drop-down list of compatible data fields for the value to place in the event field. This works for all field types. • Use an expression in the format <code>@<eventfieldName></code> to set a string type field such as Device Custom String1 with the value of a unique aggregation field. For example, if you are doing unique aggregation on source address, your value for Device Custom String 1 = <code>@sourceAddress</code>. Setting event field values for unique aggregation fields are only supported on these rule triggers: on first threshold, on every threshold, on subsequent threshold, on time unit, and on time window expiration. See procedures in "Setting or Changing Rule Thresholds" on page 528 for a description of unique aggregation fields. If the correlation event already has a value for the selected data field, that value is overridden with this rule action. Notes: • Set Event Field is the only available action for pre-persistence rules. If a pre-persistence rule calls the Set Event Field action, the modification is done to the incoming base event which has not yet persisted, instead of on the rule's correlation event. • When you edit this rule action and select new fields, these fields are added to the existing list of fields. If you want to replace the existing fields with your new fields, click the Override Fields checkbox and save the rule. • This rule action takes precedence when you are setting an event's Stage attribute. If you want to override this specific behavior, see Stage in "Annotating an Event" on page 296 for the instructions to set the override property.

Rule Actions, continued

Action	Expanded Menu Option	Description
Send Notification		Send e-mail or cell phone messages to the ESM users in the notification group when rules are triggered. Specify a notification group in the Destination Group drop-down menu, then enter the notification text in the Message box. - Click Ack Required if you want to begin an escalation chain. In this case those notified must acknowledge that they received the notification. - If you do not select Ack Required, the message is for information purposes only and is displayed on the Notifications manager's Informational tab. - For more information, see "Managing Notifications" on page 216.
Execute Command		Execute a command when the rule triggers. Select an operating system platform from the drop-down menu. - Enter the command string in the Command field. - Enter any required parameters in the Parameters field. Otherwise the command cannot execute without user intervention. Caution: If you use parameters of Date/Time with the Execute Command, you must enclose the variable name in double quotes (" "). For example, to use \$endTime as a parameter to a command to be executed on a rule action, enter the parameter as "\$endTime". - Select the Action Type: Automatically run on manager: Execute the command at the ArcSight Manager without further intervention. Run on Manager with Console confirmation: Require an operator at a Console to approve the command before it executes. Run on connector(s): Send the command to the connectors that report the events.
Execute Connector Command		Execute a SmartConnector command applicable to the device reporting the events. Select the SmartConnector to execute the command. After you select a connector, the command field is populated with the commands available for that connector. Only certain SmartConnectors can process commands beyond the basic set that all SmartConnectors support (start, stop, pause, continue, and terminate). This is similar to "Sending Control Commands to SmartConnectors" on page 195.


Rule Actions, continued

Action	Expanded Menu Option	Description
Export to External System		Send the rule and the triggering events to an external system that is integrated with ArcSight. The export file in XML format is stored in the ArcSight Manager's archive/exports directory.
Case	- Create New Case - Add to Existing Case - Create Ticket in ServiceNow® ITSM	- Create a case when the rule is triggered. See "Using a Rule to Create a Case" on page 546 for detailed steps. The rule is designated as the case creator. - Update an existing case. See "Using a Rule to Add to an Existing Case" on page 547 for detailed steps. If you configure this rule action and the specified case does not yet exist, this update action will create the case. - Create a ticket in ServiceNow. See "Using a Rule to Create a Ticket in ServiceNow® ITSM" on page 549 for detailed steps. This rule action assumes proper configuration of both the external ticketing system and the global credentials. When the rule is triggered, the correlation event is added to the case. Tip: A suggested approach to creating and updating cases based on triggered rules is to: 1. Configure an action to create a case on first event or some other threshold, set the new case's attributes, and then 2. Add to that same case when subsequent events or thresholds are triggered for that same rule.

Rule Actions, continued

Action	Expanded Menu Option	Description
Active List	Add to Active List	Add the associated events to an existing active list that you select.
	Remove from Active List	Remove the associated events from an existing active list that you select.
Notes: • Add To Active List and Remove From Active List either take no arguments (if acting on an event-bound active list) or a list of event fields (if not dealing with an event-bound active list). The values from the specified fields (those specified either by an event-bound active list or by the argument list) form an item that is added to, or removed from, the active list. Removing an item that is not present does not cause an exception. Adding an item that is already present simply increments that item's counter. You can see this counter in the Active Lists Editor. (See "Active Lists" on page 820 and "List Authoring" on page 486 for more information.) • When you are specifying fields to be added to or removed from the active list, you have the option to select local variables from the Fields tab or global variables from the Global Variables tab. • When you add a rule action to an active list, you can choose to disable audit events for the rule action. Select the Disable Audit Event check box if you want to disable audit events for the rule action. You may want to disable audit events for a rule action if you see many audit event errors for rules that cannot be removed from the active list because they have been evicted or have expired before they were removed. Note: If you disable audit events for a rule action, it will make troubleshooting active lists more difficult. • For lightweight rules, only the Active List and Session List actions are enabled. • See the Caution box in "Specifying Rule Thresholds and Aggregation" on page 528 about aggregation settings combined with rule actions that add entries to multi-mapped active lists and overlapping session lists. • See also Use of velocity expressions in rule actions involving lists .		

Rule Actions, continued

Action	Expanded Menu Option	Description
Session List	Add to Session List	Add the associated events to an existing session list that you select.
	Terminate Session List	- Add the events to the session list when a session terminates. - Terminate the oldest session. If checked, the oldest session is added to the “terminate” session list. Oldest time is based on the session's Start Time.  Caution: If your session list has a field of type Date, and that field is mapped to Manager receipt time or End time, do not use this rule action to terminate a session list entry. Instead, use either the Entry Expiration Time or TTL Days attribute for your list. See "Creating or Editing a Session List" on page 500 for complete details about session list attributes. You can also right-click the entry on a session list viewer and select Terminate Session Entry.

Rule Actions, continued

Action	Expanded Menu Option	Description
	Notes: - When you are specifying fields to be added to the session list, you have the option to select local variables from the Fields tab or global variables from the Global Variables tab. - For lightweight rules, only the Active List and Session List actions are enabled. However, lightweight rules cannot remove entries from session lists. - See the Caution box in "Specifying Rule Thresholds and Aggregation" on page 528 about aggregation settings combined with rule actions that add entries to multi-mapped active lists and overlapping session lists. - See also Use of velocity expressions in rule actions involving lists.	
Asset	Add Asset Category To Asset	Add the asset category to the associated asset. This supports the automated discovery and categorization of assets (web servers, mail servers, firewalls, and so forth) based on the type of events each asset is sending. Rules can be constructed to listen for certain types of events, and then categorize the associated asset appropriately. You also set up a condition based on which to remove the asset category from the asset , described next.
	Remove Asset Category From Asset	Remove the asset category from the associated asset. This supports automated categorization (or de-categorization) of assets along with the rule action to add an asset category (described previously) to this asset.



Note: Duplicate rule actions after a crash recovery:

If you stop ESM, it takes a checkpoint of the rules engine so that it knows what actions have been performed and where it stopped. If ESM crashes in such a way that it cannot take a checkpoint (during a power failure, for example), it returns to the last checkpoint when ESM restarts, and replays events from there. Any actions that occurred between that checkpoint and the ESM crash are therefore repeated. Repeated actions that generate audit events generate duplicate audit events.

You should investigate repeated actions that do not duplicate well. For example, if an action adds an item to an Active List, that item's counter will be incremented. If the action runs a command, it will run the command again, and so on.

Applying Rule Actions on Cases

This topic covers details on rule actions to create or update a case.

Using a Rule to Create a Case

The **Create a New Case Action** panel provides all fields necessary to set case attributes. The following example shows that creation of the case is based on the On First Event trigger for this rule:

When: On First Event

Name: Suspicious Login Attempts \$attackerAddress

Description:

Group: Select a Case group

Owner: Select a User

☐ Include Base Events in Case

More Attributes:

Initial Follow Up Final

Attributes Description Security Classification

☐ Ticket

Ticket Type	Internal
Stage	Queued
Frequency	0-0<1
Operational Impact	0-No Impact
Security Classification	1-Unclassified
Consequence Severity	0-None

OK Cancel Help

When you use a rule to create a case, you are required to give the case a name. It is possible that your organization has customized the Case resource so that additional fields are mandatory. If so, note that the rule will not check for the additional mandatory fields, and the rule **will not** fail if you do not set those mandatory fields with this rule action. You should therefore remember that after the rule creates the case, the case owner is responsible for editing and setting the other mandatory fields.



Note: After the rule action creates the case:

- The rule is identified as the creator of that case. You can view this information in the case editor's **Attributes** panel. The Created By field displays the rule's URI.
- The case's **History** panel contains text stating that the case was created by <rule_name> and also provides the rule's URI.

To apply the Create New Case rule action:

1. Required: Provide a name for the case.

You can name a case in conjunction with an existing field value from an event. For example, you want your action to create a new case called Suspicious Login Attempts based on a value in the event field, Attacker Address. For this scenario, your case name's format will be **Suspicious Login Attempts \$attackerAddress**.

2. Specify a case group and owner.
3. Optionally choose **Include Base Events in Case**. This means that when the rule action is triggered, both correlation events and the correlated base events that triggered the rule will be added to the case's Events tab.
4. Specify the case's attributes when the rule is triggered. For information on case attributes, refer to ["Creating or Editing a Case" on page 625](#).


For multi-line text fields, the value specified in this rule action will be appended to the existing values, but only if the new value being appended is unique. For example, your multi-line Attacker Address field has these IP addresses:

```
192.0.2.0
192.0.2.9
192.0.2.24
```

If the rule action is to set the field with Attacker Address = 192.0.2.9, this value will not be appended because the value is already in the field.

Using a Rule to Add to an Existing Case

The **Add To Existing Case Action** panel provides all fields necessary to change case attributes. The panel does not fetch current values of the existing case. Instead, the default selection is **Keep existing value**. Remember that you are defining the rule to keep or change certain case attributes *when* the action is triggered, whatever the case attributes are at that time.



To apply the Add To Existing Case rule action:

1. Specify a case by choosing one of two options:

- Select an existing case. Use the **Case** drop-down menu to navigate the Cases resource tree and select the case.



Caution: Make sure the case is not locked when the rule triggers; otherwise, the rule fails. However, even if the rule fails, correlation events are still added to the case because the addition of events does not alter the case resource itself.

- Select **Calculate case name dynamically** to specify a case defined in another rule action.

This setting does not require an existing case. If the case does not yet exist, the case is created when the rule is triggered.

For the dynamic case name, specify the name based on the same case name you provided in the other rule action that creates the case. An example of a dynamic case name is one that includes a variable. In the example, `GetMonth` is a variable name, and so the dynamic name entry is **Suspicious Login Attempts \$GetMonth**. If your variable name has spaces, replace the spaces with the underscore character. For example, if your variable is `Get Month`, then your case name is **Suspicious Login Attempts \$Get_Month**.

2. Optionally choose **Include Base Events in Case**. This means both correlation events and the correlated base events that triggered the rule will be added to the case's Events tab.
3. Specify the case's updated attributes when the rule is triggered, or keep the existing values. For information on case attributes, refer to ["Creating or Editing a Case" on page 625](#).

For multi-line text fields, the value specified in this rule action will be appended to the existing values, but only if the new value being appended is unique. For example, your multi-line Attacker Address field has these IP addresses:

```
192.0.2.0
192.0.2.9
192.0.2.24
```

If the rule action is to set the field with `Attacker Address = 192.0.2.9`, this value will not be appended because the value is already in the field.

Using a Rule to Create a Ticket in ServiceNow® ITSM

The **Create Ticket in ServiceNow® ITSM** panel provides all fields necessary to create a ticket. The following example shows that creation of the case is based on the On First Event trigger for this rule:

The screenshot shows a dialog box titled "Add 'Create Ticket in ServiceNow® ITSM' Action". At the top, it says "When: On First Event". Below this is a section titled "ServiceNow® ITSM". Inside this section, there are several fields: "Short Description" (with a placeholder "Type a value for Short Description"), "Caller Id" (with a placeholder "Type to start a search"), "Assignment Group" (with a placeholder "Type to start a search"), "Urgency" (set to "Low"), "Severity" (set to "1 - High"), "Impact" (set to "Low"), "Contact Type" (set to "Email"), "Assigned To" (with a placeholder "Type to start a search"), and "Category" (set to "Inquiry/Help"). Below these fields is a large text area labeled "Comments". At the bottom right of the dialog are three buttons: "OK", "Cancel", and "Help".

To apply the Create Ticket in ServiceNow® ITSM rule action:

1. Specify a short description and the appropriate Caller ID. For the description, you can either type a specific description or use an expression, such as \$name. All expressions common to the Action editor map to their expected values.
2. (Optional) Modify the remaining fields as appropriate, and then click **OK**.



Note: If you accept the defaults, then the resulting ticket will have the default settings as specified by the mapping. For example, if you leave the **Comments** field blank, the ticket will contain the event details. Any changes you make to the default settings override the associated mappings.

3. On the **Conditions** tab of the Rule Editor, specify the conditions for the rule, and click **Apply**.

Converting Rule Types

Where: Navigator > Resources > Rules

To convert a lightweight or pre-persistence rule to standard type:

1. Double-click the rule to open its Inspect/Edit panel
2. On the Attributes tab, change Rule Type to **Standard Rule**.

All features for standard rules are now available. You then add join conditions, change aggregation settings, and define actions on various types of triggers as required by the new standard rule.

To convert a standard rule to a different rule type:

Converting a standard rule to lightweight or pre-persistence rule requires that the rule must first meet the converted rule's requirements; otherwise, the rule you are converting will not be saved. See ["Rule Types" on page 508](#) for each type's features.

1. Make sure your standard rule already complies with the target rule type's requirements. For example, one of the requirements for a lightweight or pre-persistence rule states that the rule to be converted must have only one event condition. Refer to ["Rule Types" on page 508](#) for guidance.
2. In the Rules resource tree, right-click the standard rule you want to convert and choose **Edit Rule**.
3. In the Rules Editor, select the **Attributes** tab and change the Rule Type to **Lightweight Rule** or **Pre-Persistence Rule**.

The Aggregation tab is disabled.

4. In the **Conditions** tab, make sure that only one condition exists. If you see multiple conditions, disable the extra conditions and keep only one.
5. In the **Actions** tab:
 - a. Make sure that **On Every Event** is active and the other triggers are inactive.
 - b. For conversion to lightweight rules, make sure that the action is on an active or session list. For conversion to pre-persistence rules, the only allowed action is Set Event Field. Disable other actions.

The old settings of this former standard rule (aggregation thresholds, de-activated triggers, and disabled actions) will be restored when the rule is converted back to

standard type.

6. Save your converted rule.

Testing Rules

This information applies to standard rules.

You can test standard rules against copies of active channels for valid conditions logic, verify that rules are triggered by the events they are supposed to capture, and that they generate correlated events as expected.

The ArcSight Console provides two different ways of getting to tools for testing and verifying rules against events before deploying the rules in real time:

- Test a single rule from within the rule editor by clicking the **Test** button.
- Test rules and rule groups from the navigation tree with the **Verify Rules with Events** option.

These options are somewhat similar. They differ in the navigation paths to select or set up the channels, and more importantly in that from the rule editor you can test only the selected rule but from the navigation tree you can test several selected rules or rule groups.




Note: Only rules deployed in Real-time Rules act on live events and show up in a live channel when they are triggered. For more information, see ["Deploying Real-time Rules" on page 555](#).

Where: Navigator > Resources > Rules

To test a rule from the rule editor:

1. Right-click the rule and select **Edit Rule** to open the Rule editor for that rule in the Inspect/Edit panel.
2. Click **Test**.


This opens the Test Rule dialog where you can choose an existing active channel or create a new channel in which to verify the rule.



3. Select either **New Active Channel** or **Select an Active Channel** depending on whether you want to test the rule in a new or existing channel. If you need more help on setting up channels, see ["Creating or Editing an Active Channel" on page 224](#).

You can set override channel filters on either a new or existing active channel.

If you choose **Select an Active Channel** (which means you are opting to use an existing channel rather than create a new one), a browser displays the Active Channels resource tree for you to select the active channel.



4. Click **OK**.

The channel is displayed in the Viewer panel.

To show rule errors:

If rules have errors, the rule icon (⚡) on the Navigator changes to indicate the error.

In the Rules resource tree, right-click the rule-error icon and select **Show Error**. The error is described in a dialog box.

Debugging Rules

Debugging a rule creates an active channel with the filter condition defined in the rule.



Note: Inactive list conditions and variable field conditions are not supported. Debugging automatically removes these conditions and uses the remaining condition nodes to set the filter for the channel to test.

Where: Navigator > Resources > Rules

To debug a rule from the rule editor:

1. Right-click the rule and select **Edit Rule** to open the Rule editor for that rule in the Inspect/Edit panel.
2. Under the **Conditions** tab, right-click the main node of the filter, such as event1, and select **Debug in Channel**.

Verifying Rules with Events

This topic applies to standard rules. ArcSight Console provides two different ways to test or verify rules before deploying them. These options are somewhat similar. They differ in the navigation paths to select or set up the channels, and more importantly in that from the rule editor you can test only the selected rule but from the navigation tree you can test several selected rules or rule groups.

The first method is discussed in ["Testing Rules" on page 551](#). This topic explains how to test multiple rules or rule groups from the navigation tree using the **Verify Rule(s) with Events** option.

You can test rules by running them against a set of captured events for historical analysis. Now you can replay events to verify rules in existing active channels or create new channels for this purpose. Also, you can select a single rule, multiple rules, or a rule group to verify.




Tip: About Test Channels: A lightning bolt icon on a channel indicates it is a test channel created as a result of choosing **Verify Rules with Events** on a rule. Test channels cannot be re-used, even for the same rule. Remove test channels from the Active Channels folder in the Navigator.


Alternatives to Test Channels: If you would like to re-use a channel to test various rules, create a standard active channel, for example, "My Rules Test Channel" (see ["Creating or Editing an Active Channel" on page 224](#)), then send rules test results to that channel. You can re-use a standard channel as many times as you want to test rules (that is, *verify rules with events*).

To verify rules with events:

1. In the Rules resources tree, right-click an appropriate rule group or a specific rule and choose **Verify Rule(s) with Events**.
2. From the sub-menu, choose **More** or **New Active Channel**:
 - **More.** This displays the Active Channel Selector dialog. Use this dialog to navigate to the channel you want.



If you want to redefine or further narrow the stream of events in the selected channel, click the Override Channel Filter tab to add filters to it. The Override Channel Filter tab shows the conditions on the currently selected channel. You can add, remove, or modify the filters here.



Click **OK** to choose the selected channel with filter modifications (if any). The selected channel is displayed in the Viewer panel.



Note: Filters shown on rule verification channels are not designed for copying and re-use outside of these special rule testing channels. Rule verification channels show rule-triggered events and other non-correlation events in the channel, but the complete filtering logic that accomplishes this is not exposed.

Filter conditions on these channels display the original filter (if one is applied) and "Session ID > 0". The session ID statement is a simplified representation of the back-end filtering taking place in the special rule verification channel to limit this particular channel to show only new rule-triggered events.

- **New Active Channel...**

Selecting this option brings up a dialog where you can set up the parameters for the active channel that displays the rules in action. Provide a name for the new channel and set the other channel options as described in ["Creating or Editing an Active Channel" on page 224](#).

Click **OK** to create the new channel with your chosen settings. The new channel is displayed in the Viewer panel.

Unlike existing active channels, channels created as for rule verification purposes have a fixed time window (they become static) for qualifying events, and the events are those that qualify under the rules in the selected group. These active channels incorporate the conditions, aggregation characteristics, and actions defined for the rules in the selected group.



Note: Rules tested against pre-existing active channels are actually executed on copies of active channels the system automatically generates for this purpose. Rules run in verify mode do not generate real rule actions correlated with live or historical system events and, therefore, when they are triggered no real rule actions are impacting the system state. Only real-time rules or scheduled rules (set up to capture batched and other types of historical data) trigger real rule actions.

Once you have created and verified rules and are ready to deploy them on real-time events, move or copy the rules to your user folder under Real-time Rules. For more information, see ["Deploying Real-time Rules" below](#) and ["Scheduling Rules" on page 558](#).

Deploying Real-time Rules

After you have created and verified rules and you are ready to deploy them on real-time events, move or copy the rules to your user folder under Real-time Rules.

Rules that run in verify or test rule mode do not generate real rule actions correlated with live or historical system events and, therefore, when they are triggered no real rule actions are impacting the system state.

Only real-time rules show up in a live channel, generate correlation events, and trigger real rule actions.



Note: A special category of rules called scheduled rules can capture batched and other types of historical data, generate correlation events, and trigger real rule actions. These act similar to real-time rules, but are deployed differently. They are evaluated according to a schedule, and trigger off of historical/past events. See ["Scheduling Rules" on page 558](#) for more information.

Deploying a Rule

Where: Navigator > Resources > Rules

Use one of the following methods:

- Right-click a rule or a rule group (folder) and select **Deploy Realtime Rule(s)**.

The deployed rules you deploy are linked into the Real-time Rules folder (Shared/All Rules/Real-time Rules). This means that if you change something in the working copy of a rule in your user folder, those changes also take effect in the deployed rule and *vice versa*.

- Drag and drop the rules from your user folder to the Real-time Rules folder. Select **Copy**, **Link**, or **Move**.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

If a rule is already enabled () , it is deployed as enabled. If a rule has been disabled () during testing phase, it is deployed into real-time rules but remains disabled until you enable it. Rules must be both enabled and deployed in real-time rules to take effect in the live system.

If you enable or disable a deployed, linked rule in the original location it is also enabled or disabled in real-time rules and vice versa. For more information, see ["Enabling and Disabling Rules" on page 511](#).

Removing or Un-deploying a Rule

You can remove rules from the Real-time Rules folder, thereby “un-deploying” them from the live system.

To un-deploy a rule (beyond disabling it), select the rule in the Real-time Rules folder, right-click, and choose **Delete Rule** from the context menu.

Depending on whether the rule was linked, moved, or copied into the Real-time Rules folder, you get different options at this point.

- If the rule has been moved or copied into your working folder, you get an option to remove it or to cancel the operation.
- If the rule is a link to the original rule in your working folder, you get options to remove it from this group only, delete it entirely from all locations, or cancel the operation. A linked file is treated as a single entity, so edit actions taken on the file in any location affect all instances of it.

Managing Rule Groups

Rule groups are created to store similar groups or rules in a single location. Groups can be created within groups to meet enterprise needs. You can have a combination of standard and lightweight rules in the same rule group. Because you cannot schedule pre-persistence rules, keep them in the same rule group which you would not schedule.

 **Caution:** Do not exceed more than 10,000 resources in a group.

Move and copy groups and rules on the Rules resource tree with the drag and drop functionality. If you delete a group, the rules within that group are also deleted.



Note: To copy multiple resources at once, use Copy and Paste. You can drag and drop only one resource at a time.

Where: Navigator > Resources > Rules > *rule group*

To create a rule group:

1. Right-click a rule group and select **New Group**.
A New Group text field appears under the group you selected.
2. Enter the new group's name in the text field.
3. Press **Enter**.
4. Refer to ["Scheduling a Rule Group" on page 559](#) to add entries in the group's **Jobs** tab.

To rename a rule group:

1. Right-click a group and select **Rename**.
2. In the text field, enter the group's new name.
3. Press **Enter**.

To edit a rule group:

1. Right-click a group and select **Edit Group**.
2. In the Group Editor, edit the **Name** and **Description** text fields.
3. Optionally designate owners of a rule, and specify user groups that are notified of rules changes.
4. Click **OK**.
5. Refer to ["Scheduling a Rule Group" on page 559](#) to add entries in the group's **Jobs** tab.

To move or copy a rule group:

1. Navigate to a group and drag and drop it into another group.
2. Select **Move** to move the group, **Copy** to make a separate copy of the group, or **Link** to create a copy of the group that is linked to the original group.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item,

whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete a rule group:

1. Right-click a group and select **Delete Group**.
2. Click **Yes**.

Importing and Exporting Rules

To import and export rules, use the packages feature. Packages supersede the import/export facility provided in previous ESM releases and offer enhanced functionality, including version support, dependency management, and import/export capabilities. Portable ArcSight packages can automatically manage dependencies across resources and other packages. See the information on packages in ["Managing Packages" on page 730](#).

Scheduling Rules

You can schedule rules to run at a specified time interval such as hourly, daily, or monthly.

Scheduled rules are a useful alternative to real-time rules in situations where you want to deploy rules that take into account historical data along with live data, or when you simply want to control when the rules are run. The scheduled rules engine can process historical data, take real actions, and generate correlation events which are the same as those generated by the real-time rules engine.

Best practices for scheduled rules:

- Scheduling does not apply to pre-persistence rules. If an unscheduled rule group includes pre-persistence rules, the Console prevents you from scheduling that rule group. However, scheduling is at the group level. You can still schedule a parent group even if one or more of its child groups contain pre-persistence rules.
- Use future start and end dates for your rule schedules. If you use past dates, the rule will not run because the scheduler assumes that the job is finished.
- After a scheduled job has run once, it will not run again even if you change the schedule's start time and try to rerun it. Remove the old schedule and create a new one in this case.
- If you are interested only in historical data and you use a past start time and no end time, the rule will be in "catch up" mode. The rule will execute on all the data from the start to the current time, not only on historical data.

Topics include:

- ["Scenarios for Using Scheduled Rules" on the next page](#)
- ["Scheduling a Rule Group" below](#)
- ["Example of a Scheduled Rule \(Badge Swipes and Logins\) " on page 561](#)

Scheduling a Rule Group

Schedules are shared by all rules in a group. You cannot schedule a rule by itself, but you can schedule it as a group of one rule by putting it in its own folder.



Note: You cannot schedule pre-persistence rules, so ensure the rule group you need to schedule does not contain any. If the group contains pre-persistence rules, move them to another group that you do not need to schedule. See ["Rule Types" on page 508](#).


Where: Navigator > Resources > Rules > *rule group*

1. Identify the rules you want to schedule.
2. If these rules are not already in a rule group, create a rule group and link or move rules into it. For information on how to create and work with rule groups, see ["Managing Rule Groups" on page 556](#).
3. Select a rule group, right-click, and select **Schedule Rule Group**.
The rule group's editor opens at the Jobs tab.
4. Click **Add**. Give this new job a name and description.
5. In Filter Results by, select a filter for this group of rules.




Caution: Make sure the filter you select *does not* use the following filter conditions:

- InActiveList - System performance will degrade if the filter condition includes InActiveList.
- End Time - You will lose the time range if you add End Time in your filter conditions. Manager receipt time (MRT) already takes care of time range, so End Time is not needed.



You can also use the Advanced Selector button to find the filter by name. See ["Using the Advanced Selector While Editing Resources" on page 88](#) for related procedure.

- Click in the **Summary** field to set up the schedule frequency (hourly, daily, or weekly) and the schedule range (Start and End).



- Save the schedule for the rule group.

The rules are deployed according to the schedule specified in the Rule Group editor on the Jobs tab, and are triggered if the rule conditions are met.

Scenarios for Using Scheduled Rules

Batched events:

In many environments, certain types of events are not immediately available to the Manager. Instead, the events arrive in batches infrequently: once a day or once a week. Such events have different Manager receipt times and end times. Manager receipt times are current (when the batches are submitted), but the event end times are in the past, because they happened in the past. Common examples of events sent in batches are those involving physical security devices, and represent individuals gaining entry to buildings or offices by means of badge readers and card keys.

Since these events (like an employee entering an office) arrive late to the Manager, they cannot be effectively correlated with other events (like a user login) by deployed rules that use the real-time rules engine. When the real-time rules engine receives login events, it waits for 1 minute (or whatever the time window for this rule is) and then discards that login event, since the other event did not arrive within rule's time window. Consider a rule that looks for a badge swipe event and a login event within 1 minute of each other (aggregates on 1 minute). The login events are received by the Manager in real time as they occur. But the badge swipe events are collected and submitted only once a day at 10 p.m.

A real time rule would not correlate the two events because it would discard the login event before it ever gets the batched event. But if you scheduled your rule to run at midnight with the scheduled rules engine, ESM could correlate the actual end times of batched events and

login events that occur within 1 minute of each other. Scheduled rules can correlate these types of events because:

1. Rules can be scheduled to run when both the login and batched events are available within the database, and
2. Although the Manager receipt times for these events would be different, their end times are close together within the aggregation window. Correlations are based on end times of events.

Historical data:

You may want to capture and correlate other kinds of historical data other than batched events. For example, if you have observed a pattern of events over the last several weeks, decide to write rules to take actions on some of those events, and correlate not only future occurrences of them but also the past events. This is possible to do by specifying the desired date range in the filter for the schedule, for example, a filter that specifies Manager receipt time (MRT) to be between a past date range.

Optimized rule schedules:

Another scenario in which you might want to use scheduled rules is for rules that are more appropriate to run after business hours (for example, in the middle of the night). The job scheduler on rule groups lets you specify the appropriate schedule, and the rules are deployed as correlated events but are executed on off-hours.

In all such cases, scheduled rules generate correlation events and take real actions when triggered, just like deployed real-time rules.




Note: Although scheduled rules that correlate batched events work in part with historical data, these are deployed rules (not tests) that take actions as appropriate and do affect the live system.

Example of a Scheduled Rule (Badge Swipes and Logins)

This example applies to standard rules. The following shows the **conditions statements** for a rule that correlates Badge swipe events that are sent to the Manager in a batch file once per day; with login events that are sent to the Manager frequently in real-time. The example rule looks for an event with “swipe” in the name and an event with “login” in the name.

Example Scheduled Rule: Condition Statements

Rule has condition that matches badge swipe entry with login event. (Shown here on both Conditions Edit and Summary display tabs.)



This rule sets an **aggregation time window** to correlate these events at 2 minutes. This means that a login event (end time) must occur within 2 minutes of a badge swipe event (end time) in order for the rule to be triggered.

Example Scheduled Rule: Aggregation

The rule aggregates on 1 or more matching conditions within a 2 minute time window. A badge swipe and login entry must occur within 2 minutes of each other to be correlated and trigger the rule.

The screenshot shows the 'Inspect/Edit' dialog for a rule named 'Rule:Badge Entry and Logins'. The 'Aggregation' tab is selected. The configuration is as follows:

- # of Matches:** 2
- Time Frame:** 1 Minutes
- Aggregate only if these fields are unique:** (Empty list with 'Add...' and 'Remove' buttons)
- Aggregate only if these fields are identical:** (Empty list with 'Add...' and 'Remove' buttons)
- Summary:** Aggregate if at least 2 matching conditions are found within 1 Minutes

Buttons at the bottom: Test, OK, Cancel, Apply, Help.

Note that if you deploy this rule in real-time rules, the rule is not triggered to capture the events you want to correlate. Although the badge swipe events are actually occurring within 2 minutes of login events (according to event end times), the Manager Receipt Time for badge swipe events is always hours later (whenever they are submitted as batched events). In this kind of scenario, the real-time rules engine would never correlate these events because the badge swipe events (with late Manager Receipt time) would be read in so much later.

If, however, you deploy this as a scheduled rule to run on a nightly basis, the rule is triggered and capture the correlated events. This is because the scheduled rules engine is designed to correlate historical data with live data.

To configure this as a scheduled rule, create a new folder (group) for it under Rules resources in the Navigator, link or move the rule into the folder, then edit the rule group to add a scheduled job (on Jobs tab). The job schedule defines when the rule runs. Once the job schedule is applied to the rule group, the rule is deployed as a scheduled rule.

To create and test the example rule:


1. Create a rule called **Badge Entry and Logins**.
2. On the Conditions tab for this rule, set a condition to look for two events joined by AND; an event with **swipe** in the event name and an event with **login** in the event name.
3. Save the new rule.
4. Create a new rule group folder called **Badge Entry and Logins** and link or move the rule into that folder.
5. Edit the **Badge Entry and Logins** rule group to add a scheduled job for the rule of the same name.
6. Save the new rule group.

The rule is deployed after you save the rule group with the scheduled job.

For testing purposes, schedule the job to start in 5 minutes from the current time and then use the ArcSight Test Alert connector to test sending events to the Manager with end times within two minutes of each other and different Manager receipt times. (For example, to model a real-world scenario: set Manager receipt time for badge swipes to several hours later than for logins.)

Make sure that the start time of your scheduled job is earlier than the event end times on your test events (so that the scheduled job is running to capture the events). You should see the scheduled rule triggered on correlated events.

Start Time on Example Scheduled Rule is Set Earlier than End Times of Events



As a comparison, deploy the same rule in a real-time rules folder and send the test events again. Note that the same rule is not triggered by the real-time rules engine because it is not designed to correlate historical data.

In every scheduled run of a rule, only events arriving between that run and the earlier run are considered for input.

Chapter 18: Identity Correlation

Identity correlation provides the ability to model users and associate them with events. Identity correlation can be accomplished using **session lists** for some scenarios (*session correlation*) and **active lists** for others (*user or device correlation*).

You can capture and record session-related data in a user-defined *session list* where it can be used for a number of purposes in identifying and tracking users in relation to MAC addresses, IP addresses, machines, network logons, and so forth.

Also, you can use a pre-populated *active list* to find a value and then use the value (as a variable) in a rule. You can use this strategy to identify entities or objects in a variety of scenarios such as correlating various user IDs (logins, e-mail addresses, badge IDs) to unique IDs; mapping unique user IDs to user roles; and even finding the status of a machine by its host name.

Understanding Session Correlation

You can leverage ArcSight-provided resources (pre-defined [Session Lists](#) and [Rules](#)) or develop customized session lists to use for identity correlation, as described here.

How session correlation works:

Session correlation captures and records session-related data in a user-defined list, where it can be used by ESM's correlation engine to:

- Resolve event endpoints against DHCP sessions to identify which device was located at the reported IP address at the time of the event.
- Use existing maps that link MAC addresses or host names to users, if available.
- Attribute actions originating from a specific device to its owner.
- Extract and resolve user information from VPN logins, including the VPN user name and session characteristics.
- Track who accesses a given network node at a given time to trace events that originate from this device to users that were logged in at the time.

Session correlation is a three-step process that involves three or more ESM resources.



You define a session list, then create a rule to populate it. The results written to the session list can be used anywhere variables are used, such as to trigger other rules, or to populate active channels, dashboards, and reports.

The high-level steps are:

1. Create a session list (as described in ["Creating or Editing a Session List" on page 500](#)).
2. Create a rule to populate the session list (as described in ["Creating a Session List Rule" below](#)).
3. Use the session list output wherever needed (as described in ["Using the Session List Output" on page 568](#)).

See also ["Example: Using Session Lists to Correlate Session Data on User Logins" on page 569](#) for a walkthrough of creating and populating a session list with Windows session information.

Creating a Session List Rule

Make sure you have created a session list for this procedure.

Purpose: To create a rule that writes new sessions, or re-sends start times, into an existing session list.

Where: Navigator > Resources > Rules

Procedure:

1. Right-click a rule group and select **New Rule > Standard Rule**. The Rules Editor displays in the Inspect/Edit panel.
2. At the **General** tab, enter the following values:

Rule Settings for Session Lists

In this field...	...enter this
Name	Enter a name in the Rule Name text field. The Rule Name should be as descriptive as possible. It is stored in the Event Name data field and appears in the Event Name column of the grid view. The Rule Name text field is required and restricted to 25 characters.
Rule Type	Keep the selection, Standard Rule . Using a standard rule allows multiple event conditions, aggregations, and triggers. However, if you want to keep the rule simple, consider a Lightweight Rule , which is limited to acting on lists. See "Rule Types" on page 508 for more details.

Rule Settings for Session Lists, continued

In this field...	...enter this
Common: External ID, Alias	If this rule is referenced by an external system, such as a vulnerability scanner, enter the pertinent external ID information here. If not, leave these fields blank.
Description	Enter a description in the Description text field. The description should be meaningful and detailed. For example, "This rule creates an entry to the DHCP session list when a new DHCP session starts."
Assign: Owner, Notification Groups	Optional: Specify an owner for this resource. To notify other users automatically when this rule is changed, select existing users and notification groups from the drop-down menu.

3. On the **Conditions** tab, enter the conditions that indicate a session start and click **Apply**.
4. On the **Aggregation** tab, specify the event fields from the session list that you want to have displayed in the event grid when the rule is triggered by the session conditions specified in the Conditions tab. Aggregate all items you specified in your session list so that those values are populated when the event occurs.
5. On the **Actions** tab, set the trigger and the action you wish the rule to take when the conditions are met:
 - a. Select the trigger you want to apply to this rule and make sure it is activated.
On First Event is the default trigger. This determines which occurrence of the "session start" conditions will trigger the action to write the event to the session list as the start of a session. See ["Threshold Triggering Options" on page 535](#) for details on the available triggers.



Tip: You can use references to Velocity Templates as parameters for rule actions to derive values from event fields and variables. (See ["Velocity Templates" on page 1088](#).)

- b. While the trigger is selected, click **Add** to add an action. **Select Session List | Add to Session List**.
- c. In the Add Action dialog box at the Session List drop-down menu, navigate to the session list you created earlier. The parameters you set for the session list are displayed in the Session Field Mapping area.
- d. In the Session Field Mapping area at the Start Time field, select which event time stamp you want to use to record as the official start time.

Start Time	Description
End Time	The time the event ended.
Manager Receipt Time	The time the event arrived at the Manager.

- e. For the remaining fields you specified in your session list that have multiple choices, select which value you wish to use for your session list and click **OK**. You can find a description of the data fields, see ["Data Fields" on page 910](#).
6. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
7. Click **OK**. The relevant events matching this rule will populate the session list.

Using the Session List Output

After the session list has been populated by events that trigger the session list rule, the session data can be accessed anywhere variables can be used:

- Active channels
- Data monitors
- Dashboards
- Filters



Caution: Be careful about using large session lists in filters. This may severely impact system performance.

- Reports
- Rules

See ["Creating a Variable to Get Session List Data" below](#) for instructions.

Creating a Variable to Get Session List Data

For a given resource, create a variable using the `GetSessionData` function to get session timestamp data from a session list.

To create a variable to get session list data:

1. From the **Resources** tab in the **Navigator** pane, select the resource that will consume the session list data.
For a list of eligible resources, see [Using the Session List Output](#). This procedure uses Filters as an example.
2. Right-click a filter group and select **New Filter**.
3. On the **Attributes** tab, enter a name for the filter and set other attributes as required.
4. On the **Variables** tab, click **Add** and then select either **Local Variable** or **Global Variable** (depending on whether you want to share the variable across all resources).

5. In the Add Variable dialog, provide the following information and click **OK**:

In this field...	...enter this
Name	Enter a name for the variable. The name you enter appears in the <Lists> menu available from the Common Conditions Editor (CCE) . Spaces and special characters are allowed. Note: If you are creating a local variable, ensure that the name is unique across all resources. Local variables cannot share names.
Function	From the Function drop-down list, select List Functions > GetSessionData .
Arguments	From the <field name> drop-down list, select the session list that you created previously.
Preview	To preview the results, select an asset from the list of assets reporting events to ArcSight and click Calculate .

6. Perform any necessary session field mapping.
7. In the **Filters** tab conditions editor, scroll down to the bottom of the **Fields** list until you see **Variables**. Here you see the name of the variable you created earlier in this procedure.
In the **Operator** field, select an operator appropriate for the GetSessionData function for the variable you created.
In the **Condition** field, enter an appropriate value and then click **OK**.
Session lists that allow overlapping sessions take a comma-separated list of values. Session lists that do not allow session overlapping take a single value. This instructs the filter to derive its values from your session list.

Example: Using Session Lists to Correlate Session Data on User Logins

Using session lists for identity correlation is a three-step process that involves three or more ArcSight resources.

Prerequisite:

You need a set of Windows session events (user logins and logoffs) to verify the resources you create for this example.

High-level workflow:

1. Create a session list. In the example, the list will store information about Windows logins and logoffs.
2. Create a rule to populate it. The example will use two rules:

- A rule that is triggered at the start of a successful Windows login and populates the session list with the successful login event data
- A rule that is triggered when a user logs off and populates the session list with the session termination event data

The rules will be verified using the Verify Rules with Events tool to make sure that the rules are triggered and that your session list is populated appropriately with session logins and start/end times.

3. Use the session list output. The output in this example is a new report using the session list you just created as the data source. In general, the results written to the session list can be used anywhere variables are used, such as to trigger other rules, or to populate active channels, dashboards, and reports.

Step 1 - Create a Session List to Store Windows Sessions

Purpose: To create a session list that will contain Windows login sessions.

Where: Navigator > Resources > Lists > Session Lists tab

Procedure:

1. Right-click a user folder and select **New Session List**. (For more detailed help on creating session lists, see ["Creating or Editing a Session List" on page 500.](#))
2. Name the session list, and add the fields as shown.


Session List Attributes	Value
Name	Windows Login Sessions
Overlapping Entries	Disabled (leave unchecked) This example assumes that the Windows server we are monitoring does not support multiple-user logins, which is why we leave Overlapping Entries unchecked.
In MemoryCapacity (x1000)	10 (keep the default)

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725.](#)

3. Add the following three fields with names and types as shown. Set Username as the key-field:

Field Names for Session Lists	Type	Key Fields
Username	String	Enabled
NT Domain	String	
Device	String	

Example session list:



Step 2 - Create Rules to Populate the Session List with Windows Logins

Purpose: To create two rules with which to populate the session list:

- A rule that triggers on Windows session logins
- A rule that triggers when a Windows session terminates

More information:

- ["Rules Authoring" on page 507](#)
- ["Managing Rules" on page 509](#)

Where: Navigator > Resources > Rules

Right-click a rule group and select **New Rule**.



Tip: For this example, first create rules in a user folder under Rules for testing purposes. After you have created and verified rules and are ready to deploy them on real-time events, move or copy the rules to your user folder under Real-time Rules. Rules in Real-time Rules filter on live events that show up in a live channel when the deployed rules are triggered. See ["Deploying Real-time Rules" on page 555](#) for more information.

Rule 1: Triggers on Windows Session Logins

Create a rule to populate the session list. Use the following attributes, conditions, aggregation, and actions as shown below.

Attributes

On the **Attributes** tab, enter the name of the session login rule as follows.


- **Name:** Successful Windows Login

Conditions

Click the **Conditions** tab for the login rule, and enter the following conditions.

- Target User Name Is NOT NULL
- Target Nt Domain Is NOT NULL
- Device Host Name Is NOT NULL

Setting these conditions causes the rule to be triggered on any event that includes a device host name and a user name where the target is a Windows NT domain. (For more information on using the Common Conditions Editor or "CCE", see ["Common Conditions Editor \(CCE\)" on page 887](#) and ["Conditional Statements" on page 903](#).)




Aggregation

Click the **Aggregation** tab for the login rule. Under **Aggregate only if these fields are identical**, click **Add...** to bring up the Add Fields dialog. Select the following fields on which to aggregate and click OK to add them to the rule.


- Target User Name
- Target Nt Domain
- Device Host Name

Aggregation can be used to combine multiple events (as specified in the number of matches) into a single entry for the session list. But in this case (where we are aggregating events with identical fields on only a single match), we are specifying fields in the Aggregation tab for the purpose of making those same fields available in the Actions tab.



Actions

Click the **Actions** tab for the login rule.




Right-click **On Every Event** and select **Activate Trigger**. Right-click again and select **Add | Session List | Add to Session List**.

In the Session List drop-down menu on the Add dialog, select the Windows Login Sessions session list you created in the first step.


Map the fields as follows.

- Start Time: End Time
- Username: Target User Name
- NT Domain: Target Nt Domain
- Device: Device Host Name

This prompts the rule to add a login event to the Windows Login Sessions list every time a matching login event occurs.



Click **OK** on the Add to Session List dialog to add the actions to the rule. When the actions are properly configured, they are displayed under the On Every Event action as shown. Windows session logins are added to the session list on every event.



Click **OK** to save the session login rule.


Rule 2: Triggers on Termination of Windows Sessions

Create a rule to populate the session list with Windows session termination information. Define this "terminate session list" rule with the same settings as the "add to session list" rule you just created, with the following differences specific to terminating the session:


- On the **Attributes** tab, Rule Name is Windows User Logoff (instead of Login).
- On the **Conditions** tab, define the same Conditions as in the previous rule.
- On the **Aggregation** tab, aggregate on the same fields as in the previous rule.

- On the **Actions** tab, define the same actions as in the previous rule but add the actions to **Terminate Session List** instead of Add to Session List. The menu path for adding the logoff rule is **Add | Session List | Terminate Session List**.


The **Actions** tab for the logoff rule is shown below. Notice that for Windows logoffs, the rule triggers the action to add an entry to the terminate session list on every logoff event.



Here is an example of the **Attributes** tab for the logoff rule when it is completely configured.



Example of created rules on the Navigator:



Step 3 - Verify Rules

Purpose: To verify that the rules are working as expected using an active channel.

Rule	Verify Questions
Add to Session List	Is the rule triggered when a Windows login occurs? Are the values inserted into the Session List?
Terminate Session List	Is the rule triggered when a Windows logoff occurs? Is the End Time in the Session List changing according to the rule (that is, is it terminating the session for this user)?

More information:


See "Verifying Rules with Events" on page 553.

Where: Navigator > Resources > Rules and Navigator > Resources > Lists > Session Lists tab

To test rules before deploying:


1. Right-click the rule to be tested and select **Verify Rule(s) with Events**. You can create a New Active Channel to test the rules.

The following example shows the login rule triggered for several events.



2. Right-click your Windows Login Sessions list and select **Show Entries**.

The Viewer panel displays your session list entries in a channel.



The screenshot shows the ArcSight Viewer interface. At the top, there are tabs for 'Windows Login Sessions Details', 'ArcSight User Status', 'Database Performance Statistics', and 'Test My Rules'. Below the tabs, the 'Name' is set to 'Windows Login Sessions' and the 'Filter' is 'No Filter'. The main area displays a table of session data. The table has five columns: 'Username', 'N T Domain', 'Device', 'Start Time', and 'End Time'. The data shows multiple sessions for 'jarrod' on 'orangeblast' devices, followed by sessions for 'lucy' on 'johnnyquest' and 'orangeblast' devices. The 'Start Time' and 'End Time' are listed in PDT.

Username	N T Domain	Device	Start Time	End Time
jarrod		orangeblast	27 Sep 2016 13:43:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:44:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:45:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:46:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:47:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:48:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:49:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:50:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:51:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:52:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:53:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:54:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:55:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:56:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:57:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:58:19 PDT	27 Sep 2016
jarrod		orangeblast	27 Sep 2016 13:59:19 PDT	27 Sep 2016
lucy		johnnyquest	27 Sep 2016 14:00:04 PDT	27 Sep 2016
lucy		johnnyquest	27 Sep 2016 14:00:04 PDT	27 Sep 2016
lucy		orangeblast	27 Sep 2016 14:00:04 PDT	27 Sep 2016



Note: Once you have created and verified rules and are ready to deploy them on real-time events, move or copy the rules to your user folder under Real-time Rules. Only rules deployed in Real-time Rules filter on live events and show up in a live channel when they are triggered. For more information, see ["Deploying Real-time Rules" on page 555](#).

Step 4 - Use the Session List in a Report

You can leverage session lists in a variety of resources including reports. For example, you could use a rule to correlate multiple failed VPN logins over a short timeframe with a particular user entry in the session list. You then specify that if both conditions are met, add the user to an active list such as /Active Lists/Shared/All Active Lists/ArcSight System/Threat Tracking/Suspicious List.)

Purpose: To use the session list in a simple report.

Create a new report on the session list for this example. The steps are:

- Create a report
- Choose a report template
- Choose the session list as the data source for the report
- Run the report

Where: Navigator > Resources > Reports

To create a report showing the Windows logins:

1. In the Navigator, choose the **Reports** resource and click the **Templates** tab.
2. Expand the folder /Report Templates/Shared/All Report Templates/ ArcSight System/1 Table
3. Right-click **Simple Table Portrait** and select **New Report from Template**.
A Report Editor opens.
4. Provide a name for the report (for example, Windows Login Sessions).
5. Click the **Data** tab and select Session Lists for the Data Source type and the Windows Login Sessions list for the data source.
6. Click **OK**.
7. Click the **Reports** tab. The report you created is displayed under your user folder.
8. Select the new report, right-click and choose **Run Report** or **Run Report with Defaults** from the context menu.

Following is an example of an HTML version of the Windows Login Sessions report.

Report

[illegible]

For more information on creating and using reports, see ["Creating or Editing a Report"](#) on [page 393](#) and ["Running and Managing Reports"](#) on [page 446](#).

Example: Using Active Lists to Correlate Users

You can use active lists to find a value and then use value (as a variable) in a rule. You can use this strategy to identify entities or objects in a variety of scenarios; for example:

- Given that logins from the same attacker are showing up under multiple IP addresses, identify whether the attacks are coming from the same machine with different IP addresses.
- Correlate user logins (for example, into server machines) with physical building or room entry. A user's login ID is not the same as badge ID. You use an active list to map various

user identifiers (login, e-mail, badge) to a unique user ID (UUID) for each user.

- Map UUIDs to user roles.
- Find the current status (for example, server up or server down) of a given machine host name.
- Find the current status (for example, up or down) of a given SmartConnector.



Tip: The last two items can also be handled by data monitors.

This example shows how to build a rule that leverages unique user ID information from a pre-populated active list to correlate user logins on critical servers with badge swipe entries to the server room. The rule is triggered when a server user login does not have a matching badge swipe ID.

The example highlights how an active list with values can be leveraged for identity correlation. In this case, the active list collects target user IDs for the same user from different sources (e.g., user login, badge ID, e-mail address, phone number) and maps those different IDs to a unique user ID. The rule then uses the unique user ID to correlate badge swipe IDs with user login IDs.

For more about active lists, see also ["Creating or Editing an Active List" on page 486](#) and ["Using Rules to Populate an Active List" on page 494](#).

Example Overview

For this example, consider a scenario where server machines with critical data reside in a secure area. Only users in a specialized group are allowed physical access to the server room by swiping a badge on a card reader and user login permissions to the servers.

Assumption: This example assumes a policy against remote logins to the server room machines.

We want to monitor and correlate user access to the server room (badge swipes) and user logins on the server machines, and take action (e-mail notification) if our access policies are violated. Some examples of policy violations that we want to catch are:

- Cases where someone logged into a server but no badge swipe is registered. This could indicate policy violations such as remote logins or unauthorized server room entry.
- There is no matching badge swipe ID for a server console login (someone stole a user's badge to enter the server room, then logged in to the server with a different user ID).

This example assumes a pre-populated active list with values with a schema appropriate for storing information about user IDs. The active list keys off of user identifiers from various

sources (such as user login, e-mail address, phone number) and map these variants to the same unique user ID (UUID).

The UUID can then be used as a variable in a rule for correlating user login IDs with badge IDs. The example shows how to create this rule, which leverages the user information collected in the active list.

Step 1 - Build and Populate the Active List with User IDs

This example assumes that you have a pre-populated active list that maps user identifiers from various sources (badge ID, user login, e-mail, phone number) to unique user IDs (UUIDs). For the purposes of the example, we are interested in correlating badge IDs and user logins for users who log into critical servers. The active list (populated with our list of users) provides the "User Map" we need to derive each user's unique ID.

The active list definition includes the following two fields with names and types as shown. "User Identifier" is set as the key-field. This information is available in incoming events (badge swipes and user logins). Each user identifier is mapped to a UUID. Assume, for this example, that we got this mapping from IT or Human Resources departments. The UUID value is the information we'll want to extract from this list via a variable.

Field Names for Session Lists	Type	Key Fields
User Identifier	String	Enabled
UUID	String	

ActiveList : User Map

Attributes Notes

+ Add Entry

Active List

- * Name: User Map
- Optimize Data: ☐
- * Capacity (x1000): 10
- * TTL Days: 1
- * TTL Hours: 0
- * TTL Minutes: 0

(Name)
(Description)

* Data: ☐ Event-based ☒ Fields-based ☐ Key Fields

Name	Type	Sub-type
User Identifier	String	
UUID	String	

The unique user ID (UUID) that the user identifier maps to is provided here through an LDAP system or some other data source. This is the focus of the active list: to map various user IDs to the UUID. The UUID will then be used as a variable in a rule.

Populating an Active List with User Data

There are various ways to populate an active list with this kind of user information:

- Human Resources (HR) or IT database
- Identity management system
- Import from a CSV file (in the Navigator, right-click the active list and choose **Import CSV File**. See ["Importing and Exporting an Active List" on page 498.](#))
- Manually add names to the list




Tip: Note that this is a different type of task than populating an active list based on data gleaned from events (for example, ["Using Rules to Populate an Active List" on page 494.](#))

In this example, we *already have the "map" and the values we need (the unique user IDs) provided in the active list*, and we are going to feed them into a rule as a variable.

In the other example (using rules to populate the active list), we are using a rule to add items to an active list and to discover and use values as items are added to the list.

Here is an example of an active list pre-populated with user information.



User Identifier	UUID	Creation Time	Last Modified Time	Count
badge0123	SamanthaStevens	3 Feb 2017 18:21...	3 Feb 2017 18:21...	1
badgeID4156	StephanieMartinelli	3 Feb 2017 18:17...	3 Feb 2017 18:17...	1
badgeID5245	RobertJackson	28 Jan 2017 16:3...	3 Feb 2017 18:18...	1
rjackson	RobertJackson	28 Jan 2017 16:3...	3 Feb 2017 18:18...	1
rjackson@xyz.com	RobertJackson	4 Feb 2017 15:50...	4 Feb 2017 15:50...	1
samstevens	SamanthaStevens	3 Feb 2017 18:21...	3 Feb 2017 18:21...	1
samstevens@abc.com	SamanthaStevens	4 Feb 2017 15:49...	4 Feb 2017 15:49...	1
stephmartinelli	StephanieMartinelli	3 Feb 2017 18:17...	3 Feb 2017 18:17...	1
stephmartinelli@xyz.com	StephanieMartinelli	4 Feb 2017 15:50...	4 Feb 2017 15:50...	1

If you want to follow along with the example but don't have a database or spreadsheet of user information handy, you can manually add example data:

1. Build and save the User Map active list definition as described in ["Step 1 - Build and Populate the Active List with User IDs" on the previous page.](#)
2. In the Navigator, right-click the User Map active list and choose **Show Entries**.
The list is shown in the Viewer panel.
3. Click the Add Entry button  at the top right of the list to get the Active List Entry Editor.

4. Use the Active List Entry Editor to manually add user identifiers and unique user IDs. Click **Add** on the editor to add each line of data. To support the example, add at least two lines for each user. Keep the UUID the same, but the user identifiers different to illustrate the mapping.

User Identifier	UUID
badge0123	SamanthaStevens
samstevens	SamanthaStevens
badgeID5245	RobertJackson
rjackson	RobertJackson


Step 2 - Create a Rule that Uses Active List Values to Correlate User IDs

Now that we have an active list that maps various user IDs to unique user IDs (UUIDs), we can create a rule that makes use of the active list to correlate events coming from the same user with different user IDs (such as a badge swipe ID and a server login ID).

The following sections show how to define this example rule.

Attributes

On the **Attributes** tab, provide a name for the rule: **Server Room Console Login Policy Violations**



Variable

Next, we'll define a variable we can use to find unique user IDs (UUIDs) in the active list we created in the previous step ("[Step 1 - Build and Populate the Active List with User IDs](#)" on [page 582](#)).

To define a variable for finding unique UUIDs:

1. Click the **Local Variables** tab for your rule.
2. Click **Add** to begin. Provide these values for the variable definition.

Option	Specify this Value
Name	UserMap
Function	From the List category: GetActiveListValue

Option	Specify this Value
List	UserMap This is the active list we created in the previous step (" Step 1 - Build and Populate the Active List with User IDs " on page 582).
User Identifier (Active List Key field mapping)	Target User ID Use the pull-down under "Field" to select Target User ID event field. For matching events, the rule uses the value in the Target User ID field as a lookup key in the active list. For example, if the Target User ID is a login ID of "samstevens", a badge ID of "badge0123", or an e-mail address of "samstevens@example.com", all of these resolve to a unique user ID of "SamanthaStevens" in the active list mapping. The variable value passed to the rule to be evaluated in a condition would be SamanthaStevens, the UUID for any of those user identifiers.

The following example shows the variable definition on the Add Variable dialog.

- Click **OK** to save the variable.

The new variable is listed on the Local Variables tab as shown:

Conditions

We define the rule conditions so that each time a server machine login occurs, the rule conditions are evaluated. (The ServerRoomConsoleLogin condition causes this to happen.)



Tip: For more information on using the Common Conditions Editor (CCE), see "[Common Conditions Editor \(CCE\)](#)" on page 887 and "[Conditional Statements](#)" on page 903.

A comparison (Matching Event) is made between server room logins and badge swipe IDs in a 2-minute time window. The matching event uses our UserMap variable (see [Variable](#)) to get the unique ID from the active list we built in the previous step ("[Step 1 - Build and Populate the Active List with User IDs](#)" on page 582).

The rule is triggered in cases where you do not find a matching badge swipe ID for a user login.

We define the rule conditions as follows.

- The ServerRoomConsoleLogin condition finds server room machine logins via the event name and asset category. The summary of this condition is:

```
SeverRoomConsoleLogin : ( Name = Console Login AND Target Asset ID InGroup  
("/All Asset Categories/Server Room Machines/") )
```

This is the “start” condition that causes the rule conditions to be evaluated because **it is looking for server logins**.

- We define a Matching Event condition that correlates server machine logins (one type of user identifier) with badge IDs used for server room entry (another type of user identifier) based on the unique user ID (UUID) from the Active List.

We do this by using the UserMap.UUID variable we created for this purpose (see [Variable](#)).


```
Matching Event: SeverRoomConsoleLogin.UserMap.UUID = BadgeSwipe.UserMap.UUID
```

If we find a badge ID matches for all server logins, the rule is not triggered. If there is a server login with no matching badge ID within our time window, the rule is triggered.

- If someone logs in, we want to find a matching badge swipe ID for it. Since we are looking for users who logged in to servers but did not use their own badges to enter the room, we add a condition specifying that no badge swipe event (a negated Badge Swipe event) occurred for this user. So we add the event name called **BadgeSwipe** with condition Name = **Badge Swipe Event**, right-click the event name, and select **Negated**. This is to denote the event that did not occur. The summary of this condition is:

```
! BadgeSwipe : Name = Badge Swipe Event
```

The following examples show the rule conditions definition (Edit panel) and summary (Summary panel).




Matching Event Condition: Correlates server login IDs with badge swipe IDs based on unique IDs (UUIDs) gleaned from the active list. The UserMap variable is used to get UUID values from the active list.

Server Room Console Login Condition: Finds server room machine logins (via event name and asset category). This is the "start" condition that causes the rest of the rule conditions to be evaluated.

Negated Badge Swipe Event Condition: We are looking for users who logged in but did not use their own badges to enter the room. Adding this condition completes the scope of the conditions. When there is a server login, the rule correlates IDs (via the Matching Event), but triggers only if there is no matching badge swipe (this condition).

Following is an example of a Rule Conditions Summary.



Aggregation


For this example, use default aggregation settings. Aggregate on 1 match in a 2-minute timeframe.

Actions

1. Click the **Actions** tab for the rule to set up an action to take if the server room is breached.
2. Select **On First Event** (this trigger is activated by default), right-click and choose **Add > Send Notification** to bring up the Add "Send Notification" Action dialog.


3. Choose the Destination Group for the e-mail, type in a message, and click **OK** to add this action to the On First Event trigger.

For this example, we chose SOC Operators as the Destination Group. Our message is "Server room breach!".



4. Click **OK** to save the notification definition.

When the action is configured, it is displayed under the "On First Event" trigger as shown in the figure. According to this rule, a message is sent on the first trigger event; the first event in every time window that indicates a server room policy violation.



Click **OK**.

Step 2 - Create a Rule that Uses Active List Values to Correlate User IDs

Now that we have an active list that maps various user IDs to unique user IDs (UUIDs), we can create a rule that makes use of the active list to correlate events coming from the same user with different user IDs (such as a badge swipe ID and a server login ID).

The following sections show how to define this example rule.

Attributes

On the **Attributes** tab, provide a name for the rule: **Server Room Console Login Policy Violations**

The screenshot shows the 'Inspect/Edit' window for a rule named 'Server Room Console Login Policy Violations'. The 'Attributes' tab is selected, and the 'Name' field is highlighted with a red asterisk. The 'Rule Type' is 'Standard Rule'. The 'Common' section includes fields for Resource ID, External ID, Alias (Display Name), Description, Version ID, and a Deprecated checkbox. The 'Assign' section includes Owner and Notification Groups. The 'Parent Groups' section shows 'admin's Rules' with a path '/All Rules/Personal/admin's Rules/'. The 'Creation Information' section includes Created By (admin), Creation Time (2 Aug 2013 08:43:22 PDT), and Time Since Creation. The 'Last Update Information' section includes Last Updated By (admin), Last Update Time (2 Aug 2013 08:43:22 PDT), and Time Since Last Update. At the bottom, there are buttons for Test, OK, Cancel, Apply, and Help.

Variable

Next, we'll define a variable we can use to find unique user IDs (UUIDs) in the active list we created in the previous step (["Step 1 - Build and Populate the Active List with User IDs" on page 582](#)).

To define a variable for finding unique UUIDs:

1. Click the **Local Variables** tab for your rule.
2. Click **Add** to begin. Provide these values for the variable definition.

Option	Specify this Value
Name	UserMap
Function	From the List category: GetActiveListValue

Option	Specify this Value
List	UserMap This is the active list we created in the previous step (" Step 1 - Build and Populate the Active List with User IDs " on page 582).
User Identifier (Active List Key field mapping)	Target User ID Use the pull-down under "Field" to select Target User ID event field. For matching events, the rule uses the value in the Target User ID field as a lookup key in the active list. For example, if the Target User ID is a login ID of "samstevens", a badge ID of "badge0123", or an e-mail address of "samstevens@example.com", all of these resolve to a unique user ID of "SamanthaStevens" in the active list mapping. The variable value passed to the rule to be evaluated in a condition would be SamanthaStevens, the UUID for any of those user identifiers.

The following example shows the variable definition on the Add Variable dialog.

- Click **OK** to save the variable.

The new variable is listed on the Local Variables tab as shown:

Conditions

We define the rule conditions so that each time a server machine login occurs, the rule conditions are evaluated. (The ServerRoomConsoleLogin condition causes this to happen.)



Tip: For more information on using the Common Conditions Editor (CCE), see "[Common Conditions Editor \(CCE\)](#)" on page 887 and "[Conditional Statements](#)" on page 903.

A comparison (Matching Event) is made between server room logins and badge swipe IDs in a 2-minute time window. The matching event uses our UserMap variable (see [Variable](#)) to get the unique ID from the active list we built in the previous step ("[Step 1 - Build and Populate the Active List with User IDs](#)" on page 582).

The rule is triggered in cases where you do not find a matching badge swipe ID for a user login.

We define the rule conditions as follows.

- The ServerRoomConsoleLogin condition finds server room machine logins via the event name and asset category. The summary of this condition is:

```
SeverRoomConsoleLogin : ( Name = Console Login AND Target Asset ID InGroup  
("/All Asset Categories/Server Room Machines/") )
```

This is the “start” condition that causes the rule conditions to be evaluated because **it is looking for server logins**.

- We define a Matching Event condition that correlates server machine logins (one type of user identifier) with badge IDs used for server room entry (another type of user identifier) based on the unique user ID (UUID) from the Active List.

We do this by using the UserMap.UUID variable we created for this purpose (see [Variable](#)).


```
Matching Event: SeverRoomConsoleLogin.UserMap.UUID = BadgeSwipe.UserMap.UUID
```

If we find a badge ID matches for all server logins, the rule is not triggered. If there is a server login with no matching badge ID within our time window, the rule is triggered.

- If someone logs in, we want to find a matching badge swipe ID for it. Since we are looking for users who logged in to servers but did not use their own badges to enter the room, we add a condition specifying that no badge swipe event (a negated Badge Swipe event) occurred for this user. So we add the event name called **BadgeSwipe** with condition Name = **Badge Swipe Event**, right-click the event name, and select **Negated**. This is to denote the event that did not occur. The summary of this condition is:

```
! BadgeSwipe : Name = Badge Swipe Event
```

The following examples show the rule conditions definition (Edit panel) and summary (Summary panel).




Matching Event Condition: Correlates server login IDs with badge swipe IDs based on unique IDs (UUIDs) gleaned from the active list. The UserMap variable is used to get UUID values from the active list.

Server Room Console Login Condition: Finds server room machine logins (via event name and asset category). This is the "start" condition that causes the rest of the rule conditions to be evaluated.

Negated Badge Swipe Event Condition: We are looking for users who logged in but did not use their own badges to enter the room. Adding this condition completes the scope of the conditions. When there is a server login, the rule correlates IDs (via the Matching Event), but triggers only if there is no matching badge swipe (this condition).

Following is an example of a Rule Conditions Summary.




Aggregation

For this example, use default aggregation settings. Aggregate on 1 match in a 2-minute timeframe.

Actions


1. Click the **Actions** tab for the rule to set up an action to take if the server room is breached.
2. Select **On First Event** (this trigger is activated by default), right-click and choose **Add > Send Notification** to bring up the Add "Send Notification" Action dialog.
3. Choose the Destination Group for the e-mail, type in a message, and click **OK** to add this action to the On First Event trigger.

For this example, we chose SOC Operators as the Destination Group. Our message is “Server room breach!”.



4. Click **OK** to save the notification definition.

When the action is configured, it is displayed under the “On First Event” trigger as shown in the figure. According to this rule, a message is sent on the first trigger event; the first event in every time window that indicates a server room policy violation.



Click **OK**.

Step 2 - Create a Rule that Uses Active List Values to Correlate User IDs

Now that we have an active list that maps various user IDs to unique user IDs (UUIDs), we can create a rule that makes use of the active list to correlate events coming from the same user with different user IDs (such as a badge swipe ID and a server login ID).

The following sections show how to define this example rule.

Attributes

On the **Attributes** tab, provide a name for the rule: **Server Room Console Login Policy Violations**

The screenshot shows the 'Inspect/Edit' window for a rule. The 'Attributes' tab is selected. The rule name is 'Server Room Console Login Policy Violat...'. The rule type is 'Standard Rule'. The 'Common' section shows 'Resource ID' as '5jIqyP0ABABCAF3t0oo4ZAQ=='. The 'Assign' section shows 'Owner' as 'admin'. The 'Parent Groups' section shows 'admin's Rules' with a path '/All Rules/Personal/admin's Rules/'. The 'Creation Information' section shows 'Created By' as 'admin' and 'Creation Time' as '2 Aug 2013 08:43:22 PDT'. The 'Last Update Information' section shows 'Last Updated By' as 'admin' and 'Last Update Time' as '2 Aug 2013 08:43:22 PDT'. At the bottom, there are buttons for 'Test', 'OK', 'Cancel', 'Apply', and 'Help'.

Variable

Next, we'll define a variable we can use to find unique user IDs (UUIDs) in the active list we created in the previous step (["Step 1 - Build and Populate the Active List with User IDs" on page 582](#)).

To define a variable for finding unique UUIDs:

1. Click the **Local Variables** tab for your rule.
2. Click **Add** to begin. Provide these values for the variable definition.

Option	Specify this Value
Name	UserMap
Function	From the List category: GetActiveListValue

Option	Specify this Value
List	UserMap This is the active list we created in the previous step (" Step 1 - Build and Populate the Active List with User IDs " on page 582).
User Identifier (Active List Key field mapping)	Target User ID Use the pull-down under "Field" to select Target User ID event field. For matching events, the rule uses the value in the Target User ID field as a lookup key in the active list. For example, if the Target User ID is a login ID of "samstevens", a badge ID of "badge0123", or an e-mail address of "samstevens@example.com", all of these resolve to a unique user ID of "SamanthaStevens" in the active list mapping. The variable value passed to the rule to be evaluated in a condition would be SamanthaStevens, the UUID for any of those user identifiers.

The following example shows the variable definition on the Add Variable dialog.

Add Local Variable

Name: UserMap

Function: GetActiveListValue

Retrieve ActiveList value

Arguments

List: User Map

Field Mapping

For each key field, select a matching SecurityEvent field.

Name	Field	Key
User Identifier	Target User ID	☒

Preview

Set a value for each key field.

Name	Value
User Identifier	

- Click **OK** to save the variable.

The new variable is listed on the Local Variables tab as shown:

Inspect/Edit

Rule: Server Room Console Login...

Attributes Conditions Aggregation Actions Local Variables Notes

Add Edit Remove Make Global Move Up Move Down

Name	Expression
UserMap	get_activelist_value("/All Active Lists/Personal/admin's Active Lists/User Map")

Conditions

We define the rule conditions so that each time a server machine login occurs, the rule conditions are evaluated. (The ServerRoomConsoleLogin condition causes this to happen.)



Tip: For more information on using the Common Conditions Editor (CCE), see "[Common Conditions Editor \(CCE\)](#)" on page 887 and "[Conditional Statements](#)" on page 903.

A comparison (Matching Event) is made between server room logins and badge swipe IDs in a 2-minute time window. The matching event uses our UserMap variable (see [Variable](#)) to get the unique ID from the active list we built in the previous step ("[Step 1 - Build and Populate the Active List with User IDs](#)" on page 582).

The rule is triggered in cases where you do not find a matching badge swipe ID for a user login.

We define the rule conditions as follows.

- The ServerRoomConsoleLogin condition finds server room machine logins via the event name and asset category. The summary of this condition is:

```
SeverRoomConsoleLogin : ( Name = Console Login AND Target Asset ID InGroup  
("/All Asset Categories/Server Room Machines/") )
```

This is the “start” condition that causes the rule conditions to be evaluated because **it is looking for server logins**.

- We define a Matching Event condition that correlates server machine logins (one type of user identifier) with badge IDs used for server room entry (another type of user identifier) based on the unique user ID (UUID) from the Active List.

We do this by using the UserMap.UUID variable we created for this purpose (see [Variable](#)).

```
Matching Event: SeverRoomConsoleLogin.UserMap.UUID = BadgeSwipe.UserMap.UUID
```

If we find a badge ID matches for all server logins, the rule is not triggered. If there is a server login with no matching badge ID within our time window, the rule is triggered.

- If someone logs in, we want to find a matching badge swipe ID for it. Since we are looking for users who logged in to servers but did not use their own badges to enter the room, we add a condition specifying that no badge swipe event (a negated Badge Swipe event) occurred for this user. So we add the event name called **BadgeSwipe** with condition Name = **Badge Swipe Event**, right-click the event name, and select **Negated**. This is to denote the event that did not occur. The summary of this condition is:

```
! BadgeSwipe : Name = Badge Swipe Event
```

The following examples show the rule conditions definition (Edit panel) and summary (Summary panel).

Matching Event Condition: Correlates server login IDs with badge swipe IDs based on unique IDs (UUIDs) gleaned from the active list. The UserMap variable is used to get UUID values from the active list.

Server Room Console Login Condition: Finds server room machine logins (via event name and asset category). This is the "start" condition that causes the rest of the rule conditions to be evaluated.

Negated Badge Swipe Event Condition: We are looking for users who logged in but did not use their own badges to enter the room. Adding this condition completes the scope of the conditions. When there is a server login, the rule correlates IDs (via the Matching Event), but triggers only if there is no matching badge swipe (this condition).

Following is an example of a Rule Conditions Summary.

```
Matching Event:
ServerRoomConsoleLogin.Event ID = BadgeSwipe.Event ID

ServerRoomConsoleLogin :
( Name = Console Login AND Target Asset ID InGroup("/All Asset
Categories/Server Room Machines/") )

! BadgeSwipe :
Name = Badge Swipe Event
```


Aggregation

For this example, use default aggregation settings. Aggregate on 1 match in a 2-minute timeframe.

Actions


1. Click the **Actions** tab for the rule to set up an action to take if the server room is breached.
2. Select **On First Event** (this trigger is activated by default), right-click and choose **Add > Send Notification** to bring up the Add "Send Notification" Action dialog.
3. Choose the Destination Group for the e-mail, type in a message, and click **OK** to add this action to the On First Event trigger.

For this example, we chose SOC Operators as the Destination Group. Our message is “Server room breach!”.



4. Click **OK** to save the notification definition.

When the action is configured, it is displayed under the “On First Event” trigger as shown in the figure. According to this rule, a message is sent on the first trigger event; the first event in every time window that indicates a server room policy violation.



Click **OK**.

Chapter 19: Field Sets

The field sets panel provides access to resources that are used to group and extend the fields of the event and resource schema.

Field sets are named subsets of available *data fields*. Field sets can help you focus a grid view, Event Inspector, or other field array on a particular context, such as customer accounts or vulnerability.

Field sets are a shareable resource that you can manage and apply through the Field Sets resource tree in the Field Sets section of the Navigator panel. Field sets also support local and global variable data fields.

In addition to field sets based on the Security Event schema, you can create field sets based on certain resources. ArcSight supports the following types of field sets:

- **Actor field set.** An actor field set contains fields that make up the Actors resource. Actor fields are attributes to identify users and track their activity. ArcSight provides a base set of Actors fields from which you can make user-defined subsets.
- **Asset field set.** An asset field set contains fields that make up the Assets resource. Asset fields are attributes used to identify monitored assets. ArcSight provides a base set of Asset fields from which you can make user-defined subsets.
- **Case field set.** A case field set contains fields that make up the Cases resource. Case fields are attributes used to track events that have been added to cases. ArcSight provides a base set of Case fields from which you can make user-defined subsets.
- **Event field set.** An event field set is a named subset of available data fields from the ArcSight security event schema.

A base or root field set is provided for each schema type (Event, Actor, Asset, and so on) from which you can create user-defined subsets. A derived field set may inherit all or a subset of its parent's base fields, and additionally may include local or global variables not present in the parent. All field sets will have a parent (field sets created in previous versions of ArcSight will by default use the Event base field set as its parent).



Note: The ArcSightCommand Center includes a search feature, **fieldset**, that is different from the field set resource on the ArcSight Console.

The Field Sets tree presents tools for the following tasks:

Creating Field Sets

- **Who:** SOC operators, authors, and analysts concerned with traditional security-related use cases.

- **What:** A named subset of available data fields in the standard schema and the user-defined dynamic schema.
- **Why:** To narrow the fields available in the standard 400+ field event schema and the user-defined dynamic schema to make it easier to select and view fields.
- **Where:** Active channels, CCE
- **How:** See ["Creating a Field Set" below](#).

Creating Global Variables

- **Who:** SOC operators, authors, and analysts concerned with any type of use case.
- **What:** A way to derive a unique value from existing values in a data field, and the derived value itself, stored in a global variable field.
- **Why:** To make correlation, monitoring, and investigation more precise.
- **Where:** Active channels, CCE, regular field sets, other global variables
- **How:** See ["Global Variables" on page 609](#).

Creating a Field Set

Where: Navigator > Resources > Field Sets

To create a field set:

1. On the Console's Navigator panel, select **Field Sets** from the Resources drop-down menu.
2. Choose **File>New** on the Console's menu, or the **New Resource** button (), and the Field Set () command. You can also right-click a folder in the **Field Sets** resource tree and choose **New Field Set**.
3. In the Field Set Editor in the Inspect/Edit panel:
 - a. See ["Field Set Editor: Attributes Tab" on the next page](#) to enter basic attributes.
 - b. See ["Field Set Editor: Fields Tab" on the next page](#) to add data fields.
4. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
5. Click **Apply** to save the field set in the resource tree and continue editing. Click **OK** to save the set in the resource tree and close the editor.

To add a custom column to a field set:

1. Right-click a field set and select **Edit Field Set**.
2. Click **Add Custom Columns** at the bottom of the editor panel.

3. Select one or more custom columns in the Add Custom Columns dialog.

Field Set Editor: Attributes Tab

The attributes tab is where you name the field set and specify what type of field set it is.

Field	Description
Name	Enter a name for the field set that identifies what it represents.
Type	From the drop-down menu, select what type of field set it is: • Actor Field Set. Select this if the field set will contain only actor fields for use cases relating to tracking actors. • Asset Field Set. Select this if the field set will contain only asset fields for use cases relating to tracking assets. • Case Field Set. Select this if the field set will contain only case fields for use cases relating to tracking cases. • Event Field Set. Select this if the field set will contain fields from the ArcSight security event schema for event-based use cases.

For a description of what to enter in the Common fields, see ["Common Resource Attribute Fields" on page 725](#).

Field Set Editor: Fields Tab

The Fields tab is where you add the data fields to the field set.

The Field Set editor's Fields tab provides several sources from which you can select different types of fields:

- **Fields & Global Variables tab.** Use this tab to add existing user-defined fields and global variables. See ["Using the Fields & Global Variables Subtab " on the next page](#).
- **Field Sets tab.** Use this tab to add standard event and resource schema fields. This field selector is similar to those available in the CCE and active channel editors. See ["Using the Field Sets Subtab" on page 604](#).
- **Local Variables tab.** Use this tab to add one or more local variables defined on this field set's top level Local Variables tab. The added fields on this tab are available only to this particular field set. See ["Using the Local Variables Subtab" on page 604](#).

You can re-order and delete fields, and create aliases for event-based fields. For instructions, see ["Editing a Field Set" on page 607](#).




Tip: Looking for information about custom columns for field sets?

If you want to add a custom column, you need to create or define it first. For information about creating custom columns, see ["Customizing Columns" on page 250](#). For information about working with grid views, see subtopics in ["Monitoring Active Channels" on page 224](#).

After you create a custom column, you can add it to your field set using the **Add Custom Columns** button at the bottom of the Fields tab editor. For details, see ["Creating a Field Set" on page 601](#).

Using the Fields & Global Variables Subtab

The Fields & Global Variables tab enables you to select fields from a resource tree like the one presented in the Fields & Global Variables Navigator panel. Use this tab to add user-defined fields and global variables to your regular field set.




Note: The Fields & Global Variables subtab also presents regular event fields. The field selector provides a tree-level view of the standard event and resource schema fields.

In the Fields and Global Variables tab, select any existing fields or global variables you want to add to the field set. The selected field will appear in the Selected Fields panel.

For more about global variables, see ["Global Variables" on page 609](#).

Using the Field Sets Subtab

The Field Sets subtab enables you to select regular event fields that are part of a field set using a functionally organized field selector similar to that in the CCE and active channel editor. You can also use field sets in the Field Sets tab to narrow the list of fields down to those you are interested in.



You can navigate the entire event and resource schema for the fields you are interested in, or select a field set from which you want to select fields in the *Choose fields from* drop-down menu.

Using the Local Variables Subtab

If you want to add a local variable to this field set, but the local variables tab in the Field tab contains no items to select, first define the local variable in ["Field Set Editor: Local Variables Tab" on the next page](#).

1. In the Available Fields for <type of> Field Set section at the Local Variables tab, select a local variable that you have defined in [Field Set Editor: Local Variables Tab](#).
2. Select the check box for the local variable you want to add to the field set.
 - To re-order the local variables in the list, select a field and use the up /down  arrows to place it in the desired order. The variables will be evaluated in the order shown here.

- To remove the local variable from the list, select the field and click the delete button ().

Field Set Editor: Local Variables Tab

Use this top-level local variables tab to define one or more local variables that you can then add to this field set. You can create multiple chained variables and add one or more of them to the field set.

To define a local variable:

1. On the **Local Variables** tab, click **Add**.
2. In the **Name** field, give the local variable a name.



Note: Ensure that the name is unique across all resources. Local variables cannot share names.

3. From the **Function** drop-down list, select a function category and a function, and then click **OK**.
4. In the **Arguments** section, enter appropriate arguments for the function you selected in the previous step.
5. In the **Preview** section, select or enter parameters and click **Calculate** to test the function results.

For complete instructions about constructing a variable, see [Variables](#).




Tip: Fields in italics are derived from data in other fields. Derived fields appear in various places in the console, including on the Field Set Editor and the Common Conditions Editor (CCE) aggregation tabs (for example, Rules, Filters, and so forth). See also [Using Field Sets](#).

You cannot add derived fields to the field set. If you want to add a derived field to a field set, add the parent field instead.

Adding Custom Columns to the Field Set

The bottom of the Field Set Editor's Fields tab provides a button that enables you to add an existing custom columns to the field set. For more about custom columns and how to create them, see ["Customizing Columns" on page 250](#).



To add a custom column:

1. Click **Add Custom Columns**.
2. In the Add Custom Columns dialog, select an existing custom column and click **OK**.

Renaming a Column Using an Alias

The active channel bases the column title on the field's script alias, although spaces are used for better reading. For example, under the [Flex Group](#), one of the fields is Date1 with script alias = flexDate1. On the Console's event viewer or active channel, the column title for this field is Flex Date1. You can assign an alias so that the column title uses a more descriptive field that applies to your organization. After creating the alias, you can use it as a variable.

Where: Navigator > Resources > Field Sets > Fields & Global Variables tab

To assign an alias:

1. On the Fields & Global Variables tab, expand /All Fields/ArcSight System/Event Fields.
2. Right-click the field for which you want to create an alias, and select **Create Alias**.
3. In the Create Alias popup, enter the string for the alias. Keep the default group, Variables.
 - In the **Alias** field, enter the string that will be displayed as the field's column title.
 - In the **Group** field, keep the default **Variables** as the group.
 - In **Save Alias Under**, you can keep the default.

This alias is now working as a global variable.

4. Click **OK**.

You can now use the new alias as a variable for the field. Add the field to a field set, then use that field set to populate an active channel. See ["Creating or Editing an Active Channel" on page 224](#) for details, which includes applying a field set and using variables for that channel.

Editing a Field Set

Where: Navigator > Resources tab > Field Sets > Field Sets tab

1. Right-click a field set and select **Edit Field Set**.
 2. In the Field Set Editor, use the **Attributes** tab to change the field set's name.
 3. Click the **Fields** tab and use its Available Fields list to select fields to add to the list.
 - **To re-order fields:** To re-order the fields in the list, select one or more fields and use the up  or down  arrows to place fields in the desired order. Fields and variables will be displayed and evaluated in the order specified in this list.
 - **To create an alias for event-based fields:** To create an alias for a field, select the field, then click the alias button (). In the Create Alias dialog box, enter an alternate name for the field. This alias will be used to identify this field in this field set anywhere this field set is used to select or display fields, such as an active channel column heading or a CCE field selector.
- 

Note: You can create an alias for event-based fields only.
You cannot create an alias for resource-based fields, such as assets or cases. You also cannot create an alias for a field set or a global variable.
- **To delete a field from the field set:** To remove the field from the list, select the field and click the delete button ().
 4. Use the **Local Variables** tab to define variables you can add to the field set using the Local Variables tab in the Fields tab. See ["Field Set Editor: Local Variables Tab" on page 605](#).
 5. Rearrange or remove fields in the Fields to Show list.
 6. Click **Apply** to save the set in the resource tree and continue editing. Click **OK** to save the set in the resource tree and close the editor.

Sharing a Field Set

When you create a field set in the Shared folder in the Field Sets resource tree, it is available to other users who have permission for those folders. If you create one in your own folder, it is not available to other users unless you move, copy, or link it into a Shared folder.

Where: Navigator > Resources > Field Sets

1. Select the field set in your folder and drag it to the appropriate Shared folder.
2. In the Drag and Drop dialog box, choose to **Move**, **Copy**, or **Link** the resource in its new location.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

See also:

- ["Applying a Field Set to an Active Channel" on page 236](#)
- ["Sorting Events in the Active Channel" on page 238](#)

Deleting a Field Set

Where: Navigator > Resources > Field Sets

1. Right-click the field set you want to delete and select **Delete Field Set**.
2. Click **Delete** to confirm.

Resources That Use Field Sets

Use field sets in the following resources:

- **To sort active channel columns.** See ["Applying a Field Set to an Active Channel" on page 236](#).
- **To narrow the list of fields available for selecting in the CCE.** See ["Common Conditions Editor \(CCE\)" on page 887](#).
- **To define fields for a query.** See ["Building a Query" on page 320](#).

Chapter 20: Global Variables

Variables derive values from existing data fields that you can create locally in your resource to make monitoring and correlation more specific to particular scenarios.

In addition to these local variables, ESM provides a global variable resource to define a variable once, then re-use it in multiple places:

- Where you define conditions (active channels, rules, filters, data monitors, and queries)
- Where you select fields (CCE and field sets)

Because global variables are centralized and reusable, they are building blocks for advanced correlation scenarios.

Global variables are selectable in the [Common Conditions Editor \(CCE\)](#) as additional fields on the Filters or Conditions tabs, as Group By arguments for data monitors and queries, and in rule conditions and actions. You can add variables to field sets in the Field Set Editor to extend the event and resource schema with values derived from other data fields.

You can promote resource-specific local variables to global variables.

Remote Variables Processing

Variables using Group, List, and Category Model functions are evaluated on the Manager, not directly on the Console, and are referred to as “remote” variables.

These remote variables are evaluated only once on the Console for any given event or resource. Therefore, the value of the variable on the Console does not change even if the underlying data is modified that would result in a different value for the variable. New events (seen in events channels) and resources (seen in resource channels) evaluate the variable again, and you see the updated value.

Because not all variables can be calculated on the Console, there may be a delay in returning values from variables calculated remotely on the Manager.

Global Variable Dependencies

Global variables depend on a pre-defined schema. Therefore, you cannot define a global variable using the in-memory data gathered during run time in:

- Active channels
- Active and session lists
- Query viewers

- Queries
- Trends


You can display in-memory global variables as columns in active channels, but you cannot use these variables as part of a condition or filter (for example, to derive a list or query result).

Navigating to Global Variables

Where: Navigator > Resources > Field Sets > Fields & Global Variables tab

This tab displays:

- Global variable resources defined by users and in standard content
- Standard event schema fields



To view fields in the standard schema, including device custom fields:

Browse to All Fields/ArcSight System/Event Fields.

Creating or Editing a Global Variable

Caution: Do not exceed more than 10,000 resources in a group.



Caution: If you are editing a global variable, be careful about changing its name or type if the global variable is linked to other resources. Changes to its name or type could impact other resources that link to the variable. You can change its function parameters.

Where: Navigator > Resources > Field Sets > Fields & Global Variables tab

High-level steps for creating or editing a global variable:

1. If you are creating a global variable, right-click the group for the global variable you are adding and select **New Global Variable**.
If you are editing a global variable, right-click the global variable and select **Edit Field**.
2. In the Global Variable Editor in the Inspect/Edit panel, define the new or edit the existing global variable.
 - a. In the *Attributes* tab, name the global variable, specify its type, and specify the group in which to place it to help others find it in selection lists. See ["Global Variable Editor: Attributes Tab" on the next page](#).
 - b. In the *Parameters* tab, define the parameters and the functions it performs. See ["Global Variable Editor: Parameters Tab" on the next page](#).
 - c. In the *Local Variables* tab, you can optionally add a local variable, which extracts data from a field that can be used for the overall global variable. See ["Global Variable Editor: Local Variables Tab" on page 613](#).
3. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
4. Click **Apply** to apply the changes and keep the editor open; click **OK** to save changes and close the editor.

Global Variable Editor: Attributes Tab

Field	Description
Name	Enter the variable name (which must be unique in the containing group). Global variable names cannot be SQL keywords. NOTE: The value you enter here cannot be changed after the global variable is saved. If you want to change the name of the global variable after it is saved, make note of the variable attributes and re-create the variable with the desired name.
Type	From the drop-down selector, select the type of global variable you want to create: The type you choose here determines the type of fields available to this variable, and which resources can use the data derived from it. • Event Global Variable. Select this default option if you want the global variable to operate on event fields. • Asset Global Variable. Select this option if you want the global variable to operate on fields associated with assets in the network model. • Case Global Variable. Select this option if you want the global variable to operate on fields associated with cases. • Actor Global Variable. Select this option if you want the global variable to operate on fields associated with actors.
Group	From the drop-down menu, select the group in which to place your global variable. This is the group where you find the global variable in field pick lists in the CCE and Field Sets editor. The Variables group is selected by default, which means if you want to select this global variable in the pick lists, you scroll down to the Variables group. If you want to position this variable to the top group of the pick list, you select root.

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

Global Variable Editor: Parameters Tab

Use the Parameters tab to choose the category, function, and arguments necessary to supply the values.

1. On the Parameters tab's Function field, select the function that the variable uses to evaluate. First select a category, such as **Arithmetic**, then select a function from that category, for example, **Add**. Following are function categories:
 - ["Alias Functions" on page 1067](#)
 - ["Arithmetic Functions" on page 1067](#)

- ["Category Model Function" on page 1070](#)
 - ["Condition Functions" on page 1072](#)
 - ["Group Functions" on page 1073](#)
 - ["IP Address Functions" on page 1074](#)
 - ["List Functions" on page 1075](#)
 - ["String Functions" on page 1076](#)
 - ["Timestamp Functions" on page 1077](#)
 - ["Type Conversion Functions" on page 1080](#)
 - ["Value List Functions" on page 1084](#)
2. In the Arguments fields, specify the arguments (number and type parameters depending on the function), each of which may be a constant value, a field from the parent field set, or another global variable (see ["Chaining a Global Variable " on page 622](#)). For example, for the Add function which adds two numbers, your arguments will consist of the values from two specified fields to be added.
 3. For relevant functions, you can verify that the arguments you entered in the Function and Arguments fields return the values you want by entering sample parameters in the *Preview* fields.

For details about how to fill out the Function and Arguments fields, see ["Variable Definition Fields" on page 1066](#).

Global Variable Editor: Local Variables Tab

Use the **Local Variables** tab to extract a value from a field that you want to use in the overall global variable.

To extract a value:

1. Click **Add**.
2. Enter a name for the local variable, specify a function, and add arguments (number and type parameters, depending on the function).



Note: Ensure that the name is unique across resources. Local variables cannot share names.

3. Verify that the arguments you entered in the **Function** and **Arguments** fields return the values that you expect by entering sample parameters in the **Preview** fields.

For more information about the **Function** and **Arguments** fields, see [Variable Definition Fields](#).

Moving, Linking, or Deleting Global Variables

To move or link a global variable:

1. Drag and drop the global variable to another group.
2. Select **Move** or **Link**.

Move relocates the resource, leaving a single instance of it in the Navigator tree.

Link creates a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.



Note: You cannot copy global variables.

To delete a global variable:

1. In the Navigator panel, right-click the global variable and select **Delete Field**.
2. At the confirmation dialog box, click **Delete**.

If any resources depend on this variable, a warning is displayed containing the URI of the impacted resources. You can override the warning and force-delete the variable. In such cases, the dependent resources are marked invalid; you can then edit those resources and remove any orphaned references.

Promoting a Local Variable to a Global Variable

If you have an existing resource (such as a field set or rule) that contains one or more local variables that you want to re-use in other resources, it is easy to convert that variable to a global variable.

This feature is available in the following resource editors: active channels, data monitors, field sets, filters, rules, and queries.




Note: Limitation in promoting local variables for resources


- Local variables defined for data from events, actors, cases, and assets *can* be promoted to global variables.
- Local variables defined for query viewers *cannot* be promoted to global variables. Query viewers operate on queries, which have their own distinct schema for each instance. A local variable defined for a query viewer is likely only applicable to that specific query viewer.

To promote a local variable:

1. At the Local Variables tab in the resource editor, select the local variable you want to promote. This activates the Make Global button in the local variable toolbar.
2. Click the **Make Global** button.




3. On the Fields Selector, select the group to which you want to save the global variable.
4. At the prompt, click **Yes**. This removes the variable from the local variables list and makes it available to the resource as a global variable.



If you opted to replace the local variable with the global version, you can see it by viewing the condition or selected fields tab, depending on what type of resource you are working in.

5. To find the new global variable you just promoted from the global variables tree, go to **Field Sets > Fields & Global Variables**. Navigate to the group in which you saved the global variable.



The new global variable appears in the Variables hierarchy and be available to other resources.

A global variable may also chain (use as parameters) other variables that are local to a resource. A common use case is to create a complex chain of variables, and expose only the variable representing the final result as a global variable, keeping the chained intermediate variables local to their host resource.

Adding a Global Variable to a Resource

You can add a global variable to any resource such as active channels, data monitors, and field sets in which you can express a condition that uses the ["Common Conditions Editor \(CCE\)" on page 887](#). The editor for such resources would include a subtab for adding fields and global variables.


Global variables are made available to query viewers through the queries.

See also:

- ["Accessing a Global Variable Using the CCE" below](#)
- ["Adding Global Variables to an Active Channel" on page 619](#)
- ["Adding a Global Variable to a Data Monitor" on page 620](#)
- ["Adding a Global Variable to a Field Set" on page 621](#)


Accessing a Global Variable Using the CCE

Resources that use the Common Conditions Editor (CCE) provide a button that enables you to access and add a global variable to a condition statement.




Note: Only the variables whose schema type matches the given resource are displayed when on the CCE. For example, an actor-based global variable is available for an actor-based query, but not for event- or resource-based queries.

To add a global variable from the CCE:

1. In the CCE for a given resource, click the **+/-Global Variable** button .
The Global Variable Selector displays the Fields resource tree containing your selection of global variables.
2. Select one or more variables you want to add and click **OK**.
The variables are added as part of the available fields on the CCE under the Variables group.
3. On the CCE, scroll to the bottom of the available fields. You can use these variables in

condition statements for this resource.



If the resource you are working in uses a field set that contains global variables, any global variable fields included in the selected field set are also available for selection in the CCE.

Adding Global Variables to an Active Channel

When you initially create an active channel, you can only apply fields that are defined as a field set, either an existing one, or a specific field set for use only by the active channel you are defining.

Global variables can only be added to an active channel from an existing field set that contains the global variables. If an existing field set contains one or more global variables, those global variable fields become part of your active channel.

However, if you are defining the fields only for the exclusive use of the channel you are creating, the Define Grid Fields selector on the Active Channel dialog does not present global variable fields.

To viewing global variables in the Event Inspector:

When you view events in an active channel and open an event that contains a global variable field in the Event Inspector, you may need to refresh the Event Inspector view to see the global variable fields, because the Manager processes global variable data differently from regular event data.


- If the Hide Empty Rows icon  is on (so empty rows are not displayed), you may not see the global variable fields in the event inspector.

- To refresh the view, de-select, then re-select the Hide Empty Rows icon.

Adding a Global Variable to a Data Monitor

You can add a global variable to any fields-based data monitor on the attributes tab where fields are selected. Field-based data monitors include:

- Event graph
- Hierarchy Map
- Last N Events
- Last State
- Moving Average
- Statistics
- Top Value Counts (bucketized)




Where: Navigator > Dashboards > Data Monitors tab

To add a global variable to a data monitor:

1. Right-click a data monitor and select **Edit Data Monitor**).
2. In the Data Monitor editor where you can select fields, click the value field to launch the field selector. The available fields vary depending on the type of data monitor you selected.

3. In the field selector, click the **Fields & Global Variables** tab and select an available global variable. Click **OK**.



For details about how to use the data monitor editor, see ["Creating a Data Monitor" on page 266](#).

Adding a Global Variable to a Field Set

You can also add a global variable to a field set. Once you add a global variable to a field set, whenever you apply that field set in a resource, you can select the global variable directly without having to add it first.

There are five different types of field sets:

- **Actor field set.** An actor field set contains only actor-related fields. Only a global variable created using actor fields can be added to an actor field set.
- **Asset field set.** An asset field set contains only asset-related fields. Only a global variable created using asset fields can be added to an asset field set.
- **Case field set.** A case field set contains only case-related fields. Only a global variable created using case fields can be added to a case field set.
- **Event field set.** An event field set is a named subset of available data fields from the security event schema.




Note: There are also domain field sets, but you cannot create a global variable using domain fields and you cannot add a global variable to a domain field set.

Where: Navigator > Resources > Field Sets > Field Sets tab

To add a global variable to a field set:

1. Either create a new field set (right-click a group and select **New Field Set**) or edit an existing field set (right-click a field set and select **Edit Field Set**).
2. In the Field Set editor Fields tab where you can select fields, click the **Fields & Global Variables** tab and select an available global variable.



3. Click **OK**.

Chaining a Global Variable

You can “chain” variables, that is, use one variable as a function parameter for another variable. The parent (outer) variable doing the chaining can be either a local or global variable.


A variable (local or global) may be chained inside another variable only if the child (inner) variable’s return type is compatible with the outer variable's parameter type. For example, an ADD function variable can be chained inside a variable that takes a numeric parameter.



Tip: Before making one variable a function parameter of another variable, create the inner variable first, and verify that its data type is compatible with the function you want the outer variable to perform.

These steps show an example of chaining two global variables using the Global Variable Editors' Parameters tab. You can also chain a global variable in the parameters of a local variable defined in the Local Variable tab of the Global Variable editor.

1. In the Global Variable Editor's Parameters tab, select a function that matches the data type of the global variable function you want to chain. For example, if you want to perform an arithmetic function, the child (inner) variable should be a number.



2. In the Arguments section, select the inner global variable from the Global tab.
3. Verify that the arguments you entered in the Function and Arguments fields return the values you want by entering sample parameters in the *Preview* fields.

In the case of global variables that perform lookups from Active or Session Lists, the nested sub-fields (representing the list columns) are also available for selection, provided the sub-fields are the required data type.


Chapter 21: Case Management and Queries

A case contains information about an incident, usually with one or more events attached. Use cases to track, investigate, and resolve events. Where cases are similar, you can copy events directly from one case to another. You assign cases of interest to analysts, who can investigate and resolve them based on severity and enterprise policies. You can also use rules to automatically open or update a case when certain conditions are met.


You can assign cases to groups of users who receive a notification with access to the case and its associated data. Those users can take action on the assigned case and specify other actions to be taken, assign it to another user, or resolve the case.

Cases track individual or multiple related events and export event data to third-party products. Cases can stand alone or integrate with a third-party case management system.

The Case Editor has the following features:



- The case's summary is displayed at the top of the Case Editor. The example shows the top part of the panel for a case that is about to be created. The editor for an existing case has more information. It is updated as the case attribute changes.
- The icon bar provides options to display fields for setting case attributes. The default view of the Case Editor opens at the Initial view, Attributes panel.
- For existing cases, the status summary displays more information as the cases are updated:



- For existing cases, the owner is displayed below the icon bar. If there are multiple owners, the list may be hard to read. If so, resize the panel.
- Additional panels and fields are displayed by the More Options widget.

Creating or Editing a Case

This topic describes how to create or edit a case as part of an existing case group. You can have more than 1,000 cases per case group. If you have 1,000 cases or less, you can view them by expanding the case group node on the resource navigator tree. Beyond 1,000 cases in a group, create an active channel or query viewer based on a filter for case group ID to view all cases in the group. See ["Monitoring Events" on page 224](#) for details about active channels and query viewers.

About the Case Editor icons:

Case Editor Icons	Description
 Initial icon	Contains general information about the case, classified as follows: • The Attributes panel. This is the default starting view for the Case Editor. • The Follow Up panel • The Final panel
 History icon	Displays the History panel. Contains a log of tasks performed on the case, including the timestamp on when the tasks were performed and who performed the tasks. If a rule created or updated the case, the rule's name and URI would be included in the log.
 Events icon	Displays the Events panel. Use this panel to add events you want to associate with the case. A rule action can also include correlated and base events to a case. If so, the events are added to this panel.
 Attachments icon	Displays the Attachments panel. Use this panel to attach files relevant to the case, or remove them as required. Images can include query viewers, dashboards, and data monitors that are relevant to case monitoring.
 Email icon	When clicked, opens an email popup with three fields: To , Subject , and a text field for the message. Use this to send email from the case. Click Send to transmit the email.

High level steps to create or edit a case:

Where: Navigator > Resources > Cases

1. To create a case, right-click a case group and select **New Case**.

To edit a case, right-click the case on the resource navigator and select the **Lock Case** to prevent others from modifying the case while you are editing it. Right-click again and select **Edit Case**. If you forget to lock the case first, the case's editable fields are disabled. If you forgot to lock the case first, click the **Lock Case** checkbox at the bottom of the panel.



Caution: Rules can also automate making updates to existing cases based on event conditions. However, if a case is locked when the rule triggers, the rule action will fail to update the case.

For details on the rule action on cases, refer to ["Rule Actions Reference" on page 539](#).

2. Set case initial attributes. See ["Entering Case Attributes" on page 628](#) for details.
3. Enter detailed descriptions. See ["Entering Case Descriptions" on page 631](#) for details.
4. Set security classifications. See ["Entering the Case Security Classifications" on page 632](#) for details.
5. Set Initial > Follow Up information. See ["Entering Follow Up Items for the Case" on page 633](#) for details.
6. As the case goes through inspection and getting to resolution, set fields in the **Initial > Final** panel. See the following links:
 - a. ["Entering Attack Mechanism Information" on page 633](#) for details
 - b. ["Entering Attack Agent Information" on page 634](#) for details.
 - c. ["Entering Incident Information" on page 635](#) for details.
 - d. ["Entering Vulnerability Information" on page 636](#) for details.
 - e. ["Entering Miscellaneous Information" on page 636](#) for details.
7. View events associated with the case. See ["Using the Case Events Panel" on page 642](#) for details.
8. Attach files to the case. See ["Attaching a File to a Case" on page 644](#).
9. Optional: ["Using the Case's History Panel" on page 637](#).
10. Click **OK**.

More information about cases created or edited by rules:

- Rule actions can automate the creation and modification of cases. See ["Rule Actions Reference" on page 539](#) and ["Applying Rule Actions on Cases" on page 546](#).

- If a rule creates a case:
 - The rule is identified as the case creator in the case editor's **Attributes** panel. The **Created By** field displays the rule's name and URI.
 - In the **History** panel, the text states that the case was created by <rule_name> with the rule's URI.

Locking and Unlocking Cases

Every time you edit a case, lock the case. Before editing a case, verify that it is unlocked. After editing a case and saving your changes, make sure to unlock the case.

Purpose:

- To enable the fields for editing.
- To prevent others from making changes while you are working on the case.

Prerequisite:

You must be the case owner, one of the case owners, or a member of the case's owner group (s).

To lock a case:

- On the resource navigator, right-click the case resource and select **Lock Case**.
Or
- Display the Case Editor and click the **Lock case** check box at the bottom of the panel.

Locking the case enables fields for editing and prevents case co-owners from editing and overwriting your changes. When you are done with changes, unlock the case to make the it available to co-owners.

To unlock a case:

- On the resource navigator, right-click the case resource and select **Unlock Case**.
Or
- Display the Case Editor and clear the **Lock case** check box at the bottom of the panel.

To view who locked the case:

Display the Case Editor of the locked case. At the bottom of the panel, the **Locked by** field next to the check mark displays the user name.

Entering Case Attributes

When you create a case or edit an existing case, the case editor opens on the Initial panel showing Attributes fields. This panel provides the case's basic information.

Case Section

Field	Description
Name	Required field specifying name of case.
Display ID	An identification provided by an external tracking system. Automatically assigned when a new case is first saved.

Ticket Section

Field	Description
Ticket Type	Drop-down list includes Internal , Client , and Incident types.
Stage	Indicate workflow stage of ticket; default selections include Queued, Initial, Follow-Up, Final, and Closed. See also "Creating or Editing Stages" on page 300. • Queued - this is the default stage when the case is first created in ESM. Certain basic attributes are set here, such as the case name and one or more owners. • Initial - once a case is assigned, the designated assignee changes the stage to this setting. It indicates that inspections have begun on this case. Other case attributes on other panels can now be set as applicable, such as Description and Security Classification. • Follow-Up - the case is under active investigation. • Final - the case is concluding. • Closed - the investigation is over and the case is closed. No further work is needed on the case unless it is reopened (stage is changed). If you want to enter a reason for closure, see "Using the Case's History Panel" on page 637. After the case editor UI is saved and closed, then viewed again, the stage setting is updated at the top of the edit panels.
Frequency	Indicates how often reported issue occurs. Values assigned are • 0-0<1 (never or once) • 1-1<10 (less than 10 times) • 2-10>15 (10 to 15 times) • 3-15 (15 times) • 4-x>15 (more than 15)

Ticket Section, continued

Field	Description
Operational Impact	Impact of reported issue. Values assigned are • 0-No Impact • 1-No Immediate Impact • 2-Low Priority Impact • 3-High Priority Impact • 4-Immediate Impact
Security Classification	Values assigned are • 1-Unclassified • 2-Confidential • 3-Secret • 4-Top Secret After the case editor is saved and closed, then viewed again, the security classification setting is updated for the case's summary information, top of the edit panels.
Consequence Severity	Values assigned are • 0-None • 1-Insignificant • 2-Marginal • 3-Critical • 4-Catastrophic
Reason for Closure	Values assigned are • 0-False Positive • 1-True Positive - Resolved • 2-Duplicate • 3-True Positive - Other These values are placeholders for you to customize, if you want to use this field. Refer to the <i>Cases Editor UI Customization Tech Note</i>. Familiarize yourself with the entire process of UI customization. Applicable information is covered in the topic, "Customizing Field Labels," specifically the procedure "To replace a list of string options."
Category of Situation	Default is 0-None. The value assigned is a placeholder for you to customize, if you want to use this field. Refer to the <i>Cases Editor UI Customization Tech Note</i>. Familiarize yourself with the entire process of UI customization. Applicable information is covered in the topic, "Customizing Field Labels," specifically the procedure "To replace a list of string options."
Reporting level	System-calculated number calculated based on values entered in the previous fields that pertain to the particular case.
Tip: You can use entries in all case Ticket fields to generate reports so you can categorize cases based on specific case information.	

Incident Information Section

Field	Description
Detection Time	Automatically assigned based on the first event that is added to a case. Time is based on the Manager's system time. After assignment, the value does not change even if you add events or remove existing events.
Estimated Start Time	Automatically assigned based on the Manager Receipt Time (MRT) of the oldest event attached to the case, even if more recent events have been added to the case prior to this oldest event. If you remove this oldest event from the case, Estimated Start Time takes the MRT of the next oldest event in the case, and so on. If you remove all events from the case, the field will be blank.
Estimated Restore Time	This is a user-entry field to denote the date when the case is resolved. Select a timestamp from the calendar popup.

Tip: Entering values in the **Common** section is optional. See ["Common Resource Attribute Fields" on page 725](#) for information.

Assign Section

Field	Description
Owner	The ESM user designated as owner of the case. In the resource selector, expand the user's group then select the user. Make sure the group has Write privilege to cases. • The user's group is automatically added to Owner Groups after you click Apply or OK to save the change. This means the entire group also owns the case. • As you add individual users as owners, their groups will also be added as owners if those groups are not yet owners. Owners are displayed at the top of the edit panel, below the editor icons. If the list is too long, try resizing the edit panel to make the list readable. • If the individual owner linked to multiple groups, those linked groups are also automatically added to Owner Groups. • Owners are identified at the top area of the case editor. Note: If non-ESM administrators display the case editor, they will see the owner(s), but not owner groups. See also the description of the Owner Groups field (below) to know how entries in this field are automatically affected by group owner designations.
Owner Groups	The user group designated as a collective owner of the case. The group must have Write privilege to cases. You can specify groups as owners without specifying individual users. Group inheritance is not supported in case ownership. Make sure you specify a user subgroup with members. Repeat as required. Note: If non-ESM administrators display the case editor, they will not see the owner groups. The field is empty even if it has a value. See also the description for the Owner field (above) to know how entries in this field are automatically affected by individual owner designations.
Notification Groups	One or more notification groups to be notified about the case.
Note: If you see an attribute called External Ticket, this means ESM is integrated with an external case management system. Refer to Using External Case Management Systems.	




Note: Attributes under Parent Groups, Creation Information, and Last Update Information are automatically set by ESM. For descriptions, see ["Resource Attributes" on page 1032](#).

Entering Case Descriptions

1. On the Case Editor, lock the case.
2. Click the **Initial**  icon.

3. Click the **Show more** widget and select **Initial > Description**.



The fields on this panel further describe a case. Enter information as required by your business.

Case Editor Description Tab


Field	Description
Affected Services	Text field allowing entry of up to 4000 characters.
Affected Elements	Text field allowing entry of up to 4000 characters.
Estimated Impact	Text field allowing entry of up to 4000 characters.
Affected Sites	Text field allowing entry of up to 4000 characters.

After saving and closing the case editor, the Recent Comment text box at the bottom of the panel is updated. Information is also entered in the History panel.

Entering the Case Security Classifications

While the case is under investigation, use the Security Classification panel to enter security-specific features of the case.

1. On the Case Editor, lock the case.
2. Click the **Initial**  icon.
3. Click the **Show more** widget and select **Initial > Security Classification**.



1. Enter information as required by your business, based on the following guidelines:

Security Classification Fields

Field	Description
Attack Mechanism	Options include: P-Physical , O-Operational , I-Informational , and U-Unknown
Attack Agent	Options include: I-Insider , C-Collaborative , O-Outsider , and U-Unknown
Incident Source 1	Editable text.
Incident Source 2	Editable text.


Security Classification Fields, continued

Field	Description
Vulnerability	Options include: D-Design , O-Operational , E-Operational Environment , and U-Unknown
Sensitivity	Options include: U-Unclassified , C-Confidential , S-Secret , and T-Top Secret
Associated Impact	Options include: A-Availability , C-Confidentiality , I-Integrity , and U-Unknown
Action	Selections include: B-Block/Shutdown , M-Monitoring , and O-Other
Security Classification Code	Value automatically calculated and refreshed from other Security Classification field entries above.

Entering Follow Up Items for the Case

The **Follow Up** panel contains fields where you can record information as the case is undergoing active investigation.

1. On the Case Editor, lock the case.
2. Click the **Initial**  icon.
3. Click the **Show more** widget and select **Follow Up**.



Enter information as required by your business.

Case Editor Follow-Up Tab

Field	Description
Actions Taken	Text field allowing entry of up to 4000 characters.
Planned Actions	Text field allowing entry of up to 4000 characters.
Recommended Actions	Text field allowing entry of up to 4000 characters.
Followup Contact	Text field allowing entry of up to 4000 characters.

Entering Attack Mechanism Information

The fields on this panel provide final ticket resolution and reporting information for the attack mechanism associated with a case.

1. On the Case Editor, lock the case.
2. Click the **Initial** ⓘ icon.
3. Click the **Show more** widget and select **Final > Attack Mechanism**.



4. Enter information as required by your business.

Case Editor Attack Mechanism Panel

Field	Description
Attack Mechanism	Auto-populated from Security Classification tab. See "Entering the Case Security Classifications" on page 632 for possible Attack Mechanism values.
Attack Protocol	Text field allowing entry of up to 64 characters.
Attack OS	Text field allowing entry of up to 64 characters.
Attack Program	Text field allowing entry of up to 255 characters.
Attack Time	Date field. Defaults to the current time. Set by choosing values from the calendar from the drop-down options.
Attack Target	Text field allowing entry of up to 4000 characters.
Attack Service	Text field allowing entry of up to 4000 characters.
Attack Impact	Text field allowing entry of up to 4000 characters.
Final Report Action	Text field allowing entry of up to 4000 characters.

Entering Attack Agent Information

1. On the Case Editor, lock the case.
2. Click the **Initial** ⓘ icon.
3. Click the **Show more** widget and select **Final > Attack Agent**.




Fields on this panel provide ticket resolution and reporting information related to the attack agent associated with a case.

Case Editor Attack Agent Panel

Field	Description
Attack Agent	Auto-populated from Security Classification tab. See "Entering the Case Security Classifications" on page 632 for possible Attack Agent values.
Attack Location Id	Text field allowing entry of up to 255 characters.
Attack Node	Text field allowing entry of up to 255 characters.
Attack Address	Text field allowing entry of up to 255 characters.

Entering Incident Information

1. On the Case Editor, lock the case.
2. Click the **Initial**  icon.
3. Click the **Show more** widget and select **Final > Incident Information**.




The fields on this panel provide final incident information associated with a case.

Case Editor Incident Information Tab

Field	Description
Incident Source 1	Auto-populated from Security Classification tab. See "Entering the Case Security Classifications" on page 632 for possible Incident Source 1 values.
Incident Source 2	Auto-populated from Security Classification tab. See "Entering the Case Security Classifications" on page 632 for possible Incident Source 2 values.
Incident Source Address	Text field allowing entry of up to 200 characters.

Entering Vulnerability Information

1. On the Case Editor, lock the case.
2. Click the **Initial** ⓘ icon.
3. Click the **Show more** widget and select **Final > Vulnerability**.



The fields on this panel provide final ticket resolution and reporting information related to the vulnerabilities associated with a case.


Case Editor Vulnerability panel

Field	Description
Vulnerability	Auto-populated from Security Classification tab. See "Entering the Case Security Classifications" on page 632 for possible Vulnerability values.
Vulnerability Type 1	Selections include: Accidental or Intentional
Vulnerability Type 2	Selections include: EMI/RFI, Insertion of Data, Theft of Service, Unauthorized, Probes, Root Compromise, DoS Attack, User Account, Virus, Illegal Worms, Spams, Replay/Reroute, Wiretapping, Hardware/Software, Spoofing, Unknown/New,
Vulnerability Evidence	Text field allowing entry of up to 4000 characters.
Vulnerability Source	Text field allowing entry of up to 4000 characters.
Vulnerability Data	Text field allowing entry of up to 4000 characters.

Entering Miscellaneous Information

1. On the Case Editor, lock the case.
2. Click the **Initial** ⓘ icon.

3. Click the **Show more** widget and select **Final > Other**.



The fields on this panel provide miscellaneous ticket resolution and final reporting information.

Case Editor Other Panel

Field	Description
History	Selections include: Known Occurrence and Unknown
No. Occurrences	Numeric value
Last Occurrence Time	Time selected from a calendar or drop-down options.
Resistance	Selections include: High , Low , Unknown
Consequence Severity	Auto-populated from Initial Attributes panel. See "Entering Case Attributes" on page 628 .
Sensitivity	Auto-populated from Initial Security Classification panel. See "Entering the Case Security Classifications" on page 632 .
Recorded Data	Text field allowing entry of up to 4000 characters.
Inspection Results	Text field allowing entry of up to 4000 characters.
Conclusions	Text field allowing entry of up to 4000 characters.

Using the Case's History Panel

On the Case Editor, click the **History**  icon.


The fields on the History panel provide a place to track modifications to the case or enter miscellaneous case information and notes. Historical information include user-entered comments and system-generated history on case activities. Information is initially filtered alphabetically by who or what performed the action.

For example, if the case was created by a rule, the text area of the panel states that the case was created by *rule_name*. The rule's URI is also provided. Following is a sample format:

```
Case created by rule named 'rule_name'  
URI '/All Rules/rule_group/'rule_name'
```

To perform an action:

1. On the History panel, lock the case.
2. Click the **More options** widget to display the available options:



3. Refer to the following table for available options:


Options on the History Panel

Option	What You Can Do
Reason for Closure	Selects a reason for closing this case. - In the Reason for Closure popup, select one: 0- False Positive 1-True Positive - Resolved 2-Duplicate 3-True Positive - Other - Click OK to close the popup. - Click Apply to save the case. Note: This option only tracks the reason for closure. To close the case, move the Stage to Closed.
Choose a Stage	Changes the case's stage. This is the same functionality as the Stage attribute in "Entering Case Attributes" on page 628. - In the Choose a Stage popup, select one: - Queued - Initial - Follow-Up - Final - Closed - If you have added a reason for closure, you can use this in conjunction with the previous option to close the stage. - Click OK to close the popup. - Click Apply to save the case.
Attach a File	Displays the Attach file to Case popup. See "Attaching a File to a Case" on page 644 for similar steps.
Add Comment	Displays the Add Comment popup. - Enter text up to 40,000 characters. - Click Save. - Enter text up to 40,000 characters. Click Save.
Add User	Displays the Assign User resource selector. - Expand the User resource tree to select a specific user or user group. - Click OK to enter display your selection, then click OK again to close the popup. - Click Apply to save the case. See also "Owner" on page 631.

To filter historical data:

By default, the historical data lists users who performed actions on the case.

To filter actions by date, click **Date** in the filter field and select a specific date from the calendar, for example:



Click **X** to clear your selection.

Selecting **Today** displays actions that took place on the current date. This reduces the list of historical information to the date you selected.

Working with Events in Cases

This section describes the process of adding and viewing events associated with cases. The case's Events panel displays event-related information that were added to the case.

You add events to the case through any of these methods:

- Automatically through a rule. The correlation event that triggered the creation or the update of the case is automatically added to the case's Events tab. The correlated base events are also included if the rule author selected this option. See ["Including Base Events Through a Rule" on page 643](#).
- Manually from an event grid such as active channel, data monitor, or query viewer. See ["Creating or Updating a Case from Displayed Events" on the next page](#).
- Manually from another case's event channel. See the topic, ["Copying Event Details from Case to Case " on page 644](#).

The events added to a case are attached to the case itself, therefore preserving case event history over time. See also ["Collaborating on Events \(Event Annotation\)" on page 295](#).

Other related topics :

- ["Creating or Updating a Case from Displayed Events" on the next page](#)
- ["Using the Case Events Panel" on page 642](#)
- ["Viewing a Case's Events in a Channel" on page 643](#)
- ["Deleting Events from a Case" on page 644](#)

Creating or Updating a Case from Displayed Events



Caution: Events added to a case are accessible in the context of that case to any user who has permissions to view or edit the case. Users who do not have permissions on the *events* themselves have permissions to view full events *in the context of a case* to which they are assigned.

As a best practice, keep this in mind when adding events to a case and setting access control lists (ACLs) on cases. For more information on ACLs, see ["Granting or Removing Resource Permissions" on page 108](#).

You can create or update cases directly from the Viewer panel while you are monitoring suspicious events. The events you select for the case are displayed on the case's Events tab.

To add events to a case:

1. Lock the case. Display the case editor, although this is not required.
2. Open an event channel.
3. On the event grid view, select one or more events.
4. Right-click and select one of the following:
 - **Add to Case > New Case** to create a new case.
 - **Add to Case > Case in Editor** if you already have the case displayed on the Inspect/Edit panel.



Note: Here are additional options for adding events to cases.

- If there are multiple Case Editors open when you choose **Add to Case > Case in Editor**, the selected events are added to the Case Editor in focus (showing on top of the others).
 - If no Case Editors are currently open but you choose **Add to Case > Case in Editor** option anyway, a new case is created and the selected events are added to it.
- **Case > Other** to display the Case Selector popup. Navigate to the case where you want to add the events, select the case, and click **OK**.

The selected events appear in the Case Editor on the Events panel.

5. For new cases, follow additional instructions in ["Creating or Editing a Case" on page 625](#).
6. Click **OK**.

Console adds the events and displays the destination case's Events tab.

7. To view the added events on this tab, expand **Other selected Event(s)**.




Tip: Events related to a case are preserved in the case's Events panel for tracking purposes even after the retention period where the events would typically age out of the database. See ["Using the Case Events Panel" below](#). However, on the channel, the events are available based on the retention period of the Default Storage Group. For details on the retention period and how to change it, refer to the Administration section of the [ArcSight Command Center Users Guide](#). See the topic, "Storage," for information about editing the retention period of Default Storage Group.

Using the Case Events Panel

On the Case Editor, click the **Events**  icon. You are not required to lock the case when viewing information in the Events panel.

This panel provides a list of events included in a case, for example:



To view a event information associated with the case:

Select an event from the event list on the Events panel. If the event is a correlation event (indicated by the  icon), expand it to display the associated base event.

- The **Event** subpanel displays field values based on the Event Inspector. You can change the field set to use a different set of fields for the selected event.

- The **Details** subpanel provide subset information and reference Web pages to the device vendor as applicable.
- The **Annotations** subpanel lists the stages associated with the event, if any. See also ["Collaborating on Events \(Event Annotation\)" on page 295](#) and its subtopics for general information on event annotation.
- The **Payload** subpanel displays the actual data transmitted with the event, if any. See ["Payload" on page 1018](#) and ["Working with Event Payloads" on page 301](#) for more about event payloads.

More information:

- ["Working with Events in Cases" on page 640](#)
- ["Creating or Updating a Case from Displayed Events" on page 641](#)
- ["Applying Rule Actions on Cases" on page 546](#)

Viewing a Case's Events in a Channel

ESM provides two ways to view the events attached to a case:

- Right-click a case on the resource tree and select **Case Details Channel**.
- Create an active channel by following the instructions in ["Creating or Editing an Active Channel" on page 224](#), and use the filter condition, **InCase**, to identify the case to be evaluated.

Both channel types are static. This means the channel evaluates the case once, and then displays the events found in the case at that time. In an active channel using the filter with an InCase condition, the filter evaluates for a True or False condition. If events are already attached to the case at that time of evaluation (the condition evaluates to True), these events are displayed on the channel, after which further evaluation stops. Even if the active channel using the InCase filter condition is set with the Continuously Evaluate option, new events added to the case itself after the initial evaluation are not displayed.

To update (refresh) the Case Details Channel or the active channel that uses the InCase filter:

Close, then re-open the channel.

Including Base Events Through a Rule

You can define a rule action to add events to a case. By default, if a rule action creates or updates a case, the case's Events tab automatically includes the correlation event generated by

the rule. You can also configure that same rule to include base events to be part of the case's events. See ["Rules Authoring" on page 507](#) and the Case rule action listed in ["Rule Actions Reference" on page 539](#). After base events are included, you can view them on the case's Events panel.

Copying Event Details from Case to Case

Within the same ESM installation, you can copy event details from Case A (the source case) to Case B (the destination case), for situations where those two cases are similar but tracked separately. If Case A is deleted, the copied events in Case B are still available on Case B's Events panel. You can only copy event details between two existing cases.

To copy events directly from the source case:

1. Open the destination case for editing. Locking the destination case is not required.
2. Display the source case's **Events** tab, right-click the event(s), and select **Add Event(s) to Case**.
3. On the Case Selector popup, select the destination case.

Console adds the events and displays the destination case's Events tab.



Note: You can also copy events into a case from any event viewer, like an active channel, query viewer, case details channel, or data monitor. See ["Creating or Updating a Case from Displayed Events" on page 641](#).

Deleting Events from a Case

1. Lock the case and display the case's editor.
2. Go to the **Events** panel.
3. Select one or more events.
4. Right-click and select **Remove from Case**.
5. Confirm the deletion.
6. Save the case.

Attaching a File to a Case

You can attach files containing useful details about the case, so that the details continue to be available when the case is re-assigned.

Make sure that since installing the you have logged in at least twice before attaching a file. A second login solves an authentication problem that can prevent you from attaching or uploading a file.




Note: Depending on your company policies, your administrator might restrict some file types in your environment, such as .exe, .sh, or .bin.

To attach a file:

1. Open an existing case's editor and click **Lock Case** to edit it.
2. Click **Attachments** () to display the Attachments panel.

The panel displays any attachments that were previously added, if any.

3. On the Attachments panel, click the plus (+) icon.
4. In the Attach to file case dialog, select a subtab:



- Select **Local Files** to choose files from your local drive or networked drives. Optional: if you are using the Local Files subtab, change the file type in the Files of type field to expand or focus your search. Then set these fields:

Field	Description
Look in	Browse to the location of the file you want to attach.
File Name	The default is the uploaded filename, which you can change for ESM's use if you want.
Attachment Name	A descriptive name for the file. This name can differ from the actual file name, and can include spaces. If you do not provide an alternative name here, the original file name is used.
Attachment Description	Optional description of the file. Note that the description is not part of the case data. Rather, it is part of the File resource. See "Files attached to cases are stored in these places:" on the next page.
Sharing	Click Share this file in Arcsight if you want to make the file available as a shared file resource on the ArcSight Manager. Then select a Files resource group for the shared attachment to be stored.
Advanced	Mime Type - Read-only field that indicates the Multipurpose Internet Mail Extensions (MIME) type of the attached file.
	Encoding - Text encoding; for example, you could select Chinese text for internationalization requirements.

Or

- Select **ArcSight Files** to choose a file from within ESM. Expand the ESM file resource tree to choose a file resource, then click **OK**. You are done.

5. Click **OK**.

To filter file attachments:

By default, the attachments are listed by filename in alphabetical order.

To sort attachments by date, click **Date** in the filter field. The attachments are listed in chronological order.

To reverse the sorting, click .

Files attached to cases are stored in these places:

- Within the case
- In the chosen File resource group for sharing purposes, if you chose to share the file with a specific Files resource group
- In /All Files/Attachments. The attachment is automatically stored here whether or not you choose to share the attachment. See ["Managing File Resources" on page 727](#) for additional information.



Caution: About attachments:

- Attaching the same file to multiple cases means you will have the same number of file resources duplicated in /All Files/Attachments.
- Space usage may eventually increase due to attachments. To manage space usage, see ["Best Practices on Attaching Files to a Case" on page 649](#).

To detach files from the case:

1. Open an existing case's editor and click **Lock Case** to edit it.
2. Go to the **Attachments** panel and select the attachment you want to detach.
3. Click the trash icon next to the attachment (). Then confirm by clicking **OK**.

If the file is attached to only one resource (for example, just to this case), the detached file is also deleted from the All Files/Attachments resource. But if the file still attached to another resource (for example, another case), then the file resource is not deleted.

Attaching a Data Monitor, Dashboard, or Query Viewer to a Case

Purpose:

If you previously created [Dashboards](#), [Data Monitors](#), or [Query Viewers](#) pertinent to a case, you can establish links to these resources from a case for convenience. These resources are attached as images to the case.

The image file type is jpeg (.jpg).

Assumptions:

The data monitor, dashboard, or query viewer should be available for selection. It is not required that the case exists, you can use this same feature to create a case then attach the desired resource as an image.

To attach a resource as an image:

1. If you are adding to an existing case, right-click that case and lock it. Displaying the case editor is not required. If the case does not yet exist, continue.
2. Launch the data monitor, dashboard, or query viewer you want to add to a case.
3. Right-click on the data monitor, dashboard, or query viewer and choose **Export [resource type] as Image to Case**. Then continue to select from one of these options:

Option	Description
New Case	Provide a name for the image to be attached. In the Case Editor panel for a new case, set case attributes as described in "Creating or Editing a Case" on page 625 .
Case in Editor	Use this option if you have displayed an existing case's Edit panel.
Other	Displays the case selector popup from which you select the case. Make sure that case is locked.

4. In the popup, provide a name for the image. By default, the prompt uses the resource name but you can change it. Click **OK**.

Console displays the case editor's Attachments panel that shows the attached image file. This same file is also added as a file resource in /All Files/Attachments/.

Viewing a Case Attachment

After a file is attached to a case, anyone viewing the case can view details about the file and download the image. If you are only viewing an attachment, locking the case is not required.

To view attachments to a case:

1. Open a case and click the **Attachments** () icon.
2. Right-click an attachment and select **Open** to open the file or **Download** to download the file to your local system.

If you clicked **Download**:

- a. In the file browser popup, navigate to the local directory where you want to store the file.
- b. In the **File Name** field, enter the name for the file.
- c. Click **Save**.

If the case attachment was also added as a shared resource, the file is available in the ArcSight Manager. Files resource folders. To access a shared file, choose **Files** in the Navigator and browse the folders, or choose **Edit > Find Resource** from the menus, enter the file name in the **Search query** field, and click **Find**. See ["Finding Resources" on page 704](#) for more information on this utility.

Editing a Case Attachment

Within the Attachments panel, you can only attach or detach a file. If you edit an attachment, Console displays the file resource's editor. Changes to the case attachment affect both the case attachment itself and the file resource.

Right-click options for an attached image are:

Right-click option	Function
Open	Opens the image in your default image viewer.
Download	Displays a file browser so you can provide a name and directory for the image to be downloaded.
Edit	Displays the Console's edit panel for the file resource in All Files/Attachments. Note: A case attachment also exists as a file resource in All Files/Attachments.
Share	Displays the Console's File resource selector dialog so you can select another subfolder under All Files/.

To edit a case attachment:

1. Open a locked case.
2. Click the **Attachments** () icon.
3. Right-click an attachment and select **Edit**.

The Edit panel for the file resource opens. This file resource is in /All Files/Attachments/<image_name>.

4. Follow instructions in ["Working with Files" on page 728](#) to enter your edits. For example:
 - Rename the image.
 - Replace the file.

The file resource editor has additional options not related to case attributes, as documented in [Working with Files](#).

5. Save and close the file editor.

On the case's Attachments panel, the image is changed.

Best Practices on Attaching Files to a Case

ESM automatically creates a file resource for every file that you attach to a case. These file resources are stored in /All Files/Attachments. If you are creating packages of such cases, the package will also include the actual attachments. This means the attachments exist twice in the database, as file resources and as part of the package. You should therefore consider the size of the attachment files; for example, attachments of 20 Mb and above use up a lot of space. If you are creating case packages every month for backup purposes, the space usage will increase exponentially since additional file resources are created.

You can change the size of attached files with this property in the `server.properties` file:

```
persist.file.size.max
```

The value you set will be in MB. Refer to the [ESM Administrator's Guide](#) for information on editing properties files. Note that this affects all case resources, not just one.

To manage space usage, do the following:

- Avoid attaching multiple large files to a case. Keep those files in a safe and separate repository. Create a single file containing a list of pointers to the location of related files and attach that list to the case instead.
- If cases are closed, consider packaging them and exporting the packages to a separate and

safe repository. Then delete the cases from ESM. When cases are deleted, the case attachments (the file resources) are also automatically deleted.

For more information about creating packages and exporting them, see ["Managing Packages" on page 730](#).

Closing a Case

1. Display the case's editor.
2. Lock the case, if not locked.
3. On the Attributes panel, change the Stage attribute to **Closed**.
4. In the History panel, optionally select a **Reason for Closure**. Selecting a reason for closure without changing the stage does not consider the case closed.
5. Click **OK** to save and close the case editor.

The next time you redisplay the case's editor, the status at the top of the panel indicates that Stage is Closed.

Deleting a Case



Caution: Prior to deleting cases, decide if you want to preserve them after deletion. If so, add this property (or ask an administrator to add it) in the `server.properties` file before deleting any cases:

```
case.archive_ondelate.enabled=true
```

The archived deleted cases are stored as read-only snapshots for historical purposes in the Manager's `archive/cases` directory. The filename format of the archived case is

```
YYYY-MM-DD <deleted case name>.xml
```

For important details on changing properties files, see "Managing and Changing Properties File Settings" in the [ESM Administrator's Guide](#).

There are two conditions where you **cannot delete** a case (right-click option to delete the case is disabled):

- The case is locked.
- You belong to a user group that does not have permission to delete cases. This permission, `/All Permissions/ArcSight System/Case Operations/Case Delete`, is granted through a user group's ACL setting.

To delete a case:


1. In the Cases resources tree, unlock the case you want to delete, if necessary.
2. Right-click a case and select **Delete Case**.
3. Click **Yes**.



Note: If this case is linked to another, the linked case is also deleted. For related information on linked cases, see ["Moving or Copying a Case to a Group" on the next page](#).

Granting Permission to Delete Cases

By default, new user groups added under Custom User Groups are **not** allowed to delete cases. The ability to delete cases is controlled by the permission, /All Permissions/ArcSight System/Case Operations/Case Delete, set in the group's ACL Editor on the Operations tab.



A user can belong to multiple groups. The permission to delete cases takes precedence. If at least one of those groups have permission to delete cases, then the user will have the ability to do so.

To grant or remove permission to delete cases:

1. Right-click the user group and select **ACL Editor**.
2. Follow the instructions in ["Granting or Removing Operations Permissions" on page 109](#) to grant or remove the /All Permissions/ArcSight System/Case Operations/Case Delete permission on the ACL Editor's Operations tab.
3. Go to the Resources tab:
 - a. Follow the instructions in ["Granting or Removing Resource Permissions" on page 108](#) to add the /All Cases/All Cases to the list of resources.
 - b. Check the **R** and **W** boxes to grant read and write access.

Moving or Copying a Case to a Group

Where: Navigator > Resources > Cases

Procedure:

1. Drag and drop the case into another group.
2. Select **Move** to move the case, **Copy** to make a separate copy of the case, or **Link** to create a copy of the case that is linked to the original case.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

Finding Cases

Purpose: To find an ArcSight case by specifying its reference ID or from a case grid view.

To edit a case by reference ID:

1. Right-click a group in the Cases resource tree and select **Edit Case by ID**.
2. Enter the ID string in the dialog box and click **OK** to display the case in the Case Editor.

To go to a case from the case channel view:

1. Right-click a case in the case channel view and select **Find Case in Navigator**.
2. Look for the highlighted item in the Navigator panel's Cases resource tree.

Viewing a Case's Internal Audit Events

Requirement:

To view audit events, you must belong to a group that has the following enforced filter:

/All Filters/ArcSight Administration/ESM/System Health/Events/Audit/ArcSight Audit Events

This event filter is added to the group's ACL Editor settings on the Events tab and applies to audit events for all resources. Refer to ["Adding or Removing Enforced Filters" on page 113](#) for instructions.

To view internal audit events on cases:

1. Through an active channel, query viewer, or data monitor using the proper filter, you can view details on modifications done to cases.
Refer to ["Monitoring Events" on page 224](#) for the different ways to create an event viewer.
2. Use the following fields in your event viewer to capture audit events related to the case:

Fields for Case Internal Audit Events

Field or Column	Displayed Information
Name	Event name, for example, Case updated or Note inserted.
File Name	The case's name
Target User Name	The user who made the change
Device Custom String3 (Label: AttachedTo)	For the "Note inserted" event, the resource type, Case, to which the note is attached
Device CustomString4 (Label: AttachedToID)	For the "Note inserted" event, the resource ID of the case to which the note is attached
Device CustomString5 (Label: AttachedToName)	For the "Note inserted" event, the display name of the case to which the note is attached



Tip: After the event viewer retrieves the data:

1. Copy the resource ID of interest displayed in Device Custom String4.AttachedToID).
2. Paste the ID into the Console's search box and go to the actual case resource on the Navigator panel.

Managing Case Groups

Purpose: To keep similar groups or cases in a single node on the Navigator.



Caution: Do not exceed more than 10,000 resources in a group.

Where: Navigator > Resources > Cases

To create a case group:

1. Right-click a group and select **New Group**.
2. In the **Name** text field, enter a name.

To rename a case group:

1. Right-click a group and select **Rename**.
2. In the **Name** text field, rename the group.



Caution: Verify if this case group is referenced by other resources. For example, if a rule action is triggered to update a case in this group, make sure to update the URI to the renamed group.

To edit a case group:

1. Right-click a group and select **Edit Group**.
2. In the Group Editor, edit the **Name** and **Description** text field.

To move or copy a case group:

1. Navigate to a group and drag and drop it into another group.
2. Select **Move** to move the group, **Copy** to make a separate copy of the group, or **Link** to create a copy of the group that is linked to the original group.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete a case group:

1. Right-click a group and select **Delete Group**.
2. Click **Yes** to confirm.

Deleting a case group means all cases under the group are deleted. If you delete a linked group, the original group is also deleted.

Viewing Group Cases in a Grid View

Prerequisite:

At a minimum, you should have Read permission on cases to view case information. If you want to edit a case from here, then you should have Write permission.

Do one of the following:

- From the **View** menu, select **My group's cases**.
- If you belong to the Administrators' group: from the View menu,
 - Select **My group's cases** to see cases owned by your group, or
 - Select **A group's cases** then select a group from the Users Selector popup.
- In the Cases resource tree in the Navigator panel, right-click a case group and select **Show Cases**.

The cases are displayed on a channel.

On a case channel:

- Click any case in the grid to work with the case individually.
- You can also:
 - Right-click any column heading to get a menu of column configuration options.
 - Right-click any individual case's fields to get a menu of case handling options, described as follows.

Option	Description
Edit Case	Open a case in the Inspect/Edit panel for editing.
Delete Case	Delete the selected case.
Case Detail Channel	Display the base events added to the case.
Export to external system	Export the case to an external tracking system.
Edit Case by ID	Find a case by its Display ID value.
Find Case in Navigator	Expand the group where this case belongs and highlight the case.
Lock Case	Prevent other users from making changes to the case.

Running Case Queries

Where: Navigator > Resources > Cases

You can create a specific case search group based on a the value of a particular case attribute. The search group is updated with all cases matching the condition.

To set up a case search group:

1. Right-click a case group and select **New Search Group**.
2. Give the search group a suitable name.
3. On the Conditions tab, add a case condition. For example, if you want this search group to contain all cases owned by **CaseAdmin**, the condition would say
`Owner = CaseAdmin`

The search group is populated with all cases assigned to CaseAdmin. The search group is updated as CaseAdmin's assignments change.

Creating a Report from a Case

You can run a simple case report using one of the following ways:

- If your case group has up to 1000 cases, you can view all of its cases on the resource navigator tree from which you choose the case and run the report for it.
- If your case group consists of more than 1000 cases, create an active channel based on a filter for case group ID (see ["Creating or Editing an Active Channel" on page 224](#) for additional details). On the channel, select a case and run the report for it.

Run one case report at a time.

ArcSight provides the convenience of a simple case report template that includes its underlying query that you can use as is. The output is a two-column report that is rotated so that columns (composed of the case attributes) are turned into rows, therefore making the report readable. You can also customize the template by adding case fields.

Procedures:

- ["Running Case Reports and Setting Default Parameters" on the next page](#)
- ["Customizing the Case Report" on page 659](#)


Running Case Reports and Setting Default Parameters

Where: Navigator > Resources > Cases or a case channel

To run a case report and set default parameters:

1. Right-click the case and select **Run Case Report**.

The Case Report Parameters is a simplified dialog derived from the regular report resource's parameters dialog box:



2. Set the default parameters as follows:

Report Parameter Values

Common Parameters	Description
Report Format	From the drop-down menu, select one of the following report output formats: • pdf - Outputs the report as an Adobe PDF file. • xls - Generates a Microsoft Excel file for tables and charts. Note: XLS reports you run with <i>Microsoft Excel 2002</i> might have page break format problems (misalignments, column spillover) due to default page size settings in Excel. To correct this problem, open the resulting XLS report in Excel, choose File > Page Setup from the menus, change the paper size to Letter (instead of Legal), and click OK to save your changes. The report has the appropriate page break formatting. <i>This problem does not occur in newer versions of Microsoft Excel.</i> Note: XLS report formats display speedometer charts as pie charts. This is a known limitation in Microsoft Excel. • rtf - Produces a rich-text format document. • csv - Creates tabular data as a list of comma-separated values. Note: Reports generated in CSV format are not the full equivalent of exports to other formats like PDF or HTML. CSV format is useful for loading report data into a spreadsheet for further manipulation. Since CSV is meant to contain tabular data, only the table data of a report is normally useful. Therefore, ArcSight exports only the table data portion of a report to CSV format, ignoring any other report information such as charts or text, including report titles. • html - Generates the report in a Web page displayed by the default web browser. Note: Running a report with a huge row limit (10,000 rows) in HTML format may cause the browser to be unresponsive. In such cases, run the report in PDF. Your selection affects your choice for e-mail formats. See description for Email Format, another report parameter, in this table. The output file is created at report runtime and is stored in the Archived Report Group corresponding to the report's group. See the Archives tab on the report's edit panel.
Page Size	From the drop-down menu, select a paper size.
Run as User	This option is not available for case reports. Run the report as a particular user. From the drop-down menu, select the user name by which you would like to run the report. For example, this option would allow an administrator for a Managed Security Service Provider (MSSP) to run report for a customer. The administrator would need write permissions to the user.

Report Parameter Values, continued

Common Parameters	Description
Email to	You can have the report sent as e-mail to one or more ArcSight users. From the drop-down menu, select the Console users to whom the report should be e-mailed. The selection list is read from the Users resource). The recipient will only see his or her user name in the To field even if there are multiple recipients for this report. Handling empty reports: If you are emailing reports, empty reports will also be sent. This is determined by the server property, <code>report.scheduler.notify_empty_reports</code>, which is set to true. If you do not want empty reports to be sent, add the property to the <code>server.properties</code> file and change the setting to false. Follow the instructions in the <i>ESM Administrator's Guide</i> on how to edit this file. The details are in the guide's topic, "Managing and Changing Properties File Settings."
Email addresses	Send the report to one or more comma-separated or semicolon-separated e-mail addresses. This option does not require the recipient to be an ArcSight user. Note: The recipient will only see his or her e-mail address in the To field even if there are multiple recipients for this report.
Email Format	Specify how the report is to be accessed by the recipient. • Select Send URL if you want to point users to the report. Use this option if the report is large and is saved (archived) to a network-accessible location You can provide URLs for all report formats: PDF, XLS, RTF, CSV, and HTML. • Select Attach Report if you want to send the report directly to the user's e-mail box. You can only attach PDF, XLS, RTF, and CSV report formats. • Select Attach Compressed Report if you want the PDF, XLS, RTF, or CSV report to be compressed (zipped) first before mailing. Note: The notification automatically uses the URL if you select an email format for an unsupported report format, or if the attachment is too large to be sent. You can adjust the maximum attachment size up to 100 MB with the property, <code>report.upload.maxFileSize</code>. Set this property to the size you want in the <code>server.properties</code> file. Refer to the <i>ESM Administrator's Guide</i> for procedures to edit properties files.
Email Subject	Specify the subject on the notification. Defaults to the report's Name attribute (denoted by <code>\$ReportName</code>). If you want to use a customized subject, type the text either in addition to the default or replace the default text entirely.

3. Click **OK**.
4. Respond to the prompt to open the report for viewing or to save it to a directory of your choice.

Customizing the Case Report

Customizing the case report gives you the flexibility of adding and removing case fields.

The basic case report, **Selected Case Report**, and its underlying query, **Selected Case Query**, are part of the ArcSight standard content under /All Reports/ArcSight System/Core/Selected Case Report and /All Queries/Arcsight System/Core/Selected Case Report/Selected Case Query. This case report, at a minimum, has a name and a query on a case ID. You can copy this report as a template and change the copy by adding or removing fields in the report.

By default, Selected Case Query includes the following fields for Selected Case Report:

Fields in Selected Case Report

Name	Alias	Create Time	Creator
Description	Modification Time	Owner	Consequence Severity
Operational Impact	Security Classification	Stage	Ticket Type
Reporting Level	Frequency	Detection Time	Estimated Restore Time
Estimated Start Time	Incident Source 1	Incident Source 2	Incident Source Address
Affected Elements	Affected Services	Affected Sites	Estimated Impact
Action	Associated Impact	Attack Agent	Attack Mechanism
Security Classification Code	Sensitivity	Vulnerability	Actions Taken
Followup Contact	Planned Actions	Recommended Actions	

If you want to change fields in the case report, copy both Selected Case Report and Selected Case Query into a resource group of your choice (as shown in the procedures), then modify the copies.

Follow all procedures in this topic according to the sequence:

- ["Customize Selected Case Query" below](#)
- ["Customize Selected Case Report" on the next page](#)
- ["Add a Server Property for the New Report URI" on page 662](#)

Customize Selected Case Query

This topic explains how to add fields (columns) in the report query so that they are included in the case report. In this example, we will add the following fields:

Attack OS
Attack Impact
Attack Target

To customize Selected Case Query:

1. Copy Queries/Shared/All Queries/ArcSight System/Core/Selected Case Report/Selected Case Query into a custom query group of your choice and rename it, for example:

Queries/Shared/All Queries/Public/CW3 Selected Case Query

Keeping this copy in the Public node makes the query accessible to all.



Caution: You can change the copied query's name, but keep the **Query On = Case** setting.

2. On the copied query's **Fields** tab, click **Add 'SELECT' columns** to add fields, for example: Attack OS, Attack Impact, Attack Target, and more as required. Change the sort order as desired.



Note:

The Notes field is not available for this query.

If you need additional help on creating or editing queries, see ["Building a Query" on page 320](#).

Next go to ["Customize Selected Case Report" below](#).

Customize Selected Case Report

This topic explains how to customize the basic case report by using the customized query (["Customize Selected Case Query" on the previous page](#)). For best results, use the Simple Table Portrait for your customized report. Use Simple Table Landscape if you want a wider format. If you use formats other than Table, your report readers might have usability issues.

To customize Selected Case Report:

1. Copy Reports/All Reports/ArcSight System/Core/Selected Case Report into a custom report group of your choice and rename it, for example:

Reports/Shared/All Reports/Public/CW3 Selected Case Report

2. Right-click the copied report and select **Edit Report**.
3. Go to the **Data** tab:
 - a. For the **Data Source** field, browse to the copied query (**CW3 Selected Case Query** in our example) and select it as the data source for the report.
The added fields from the modified query appear in the Available Columns table. Scroll horizontally to view all available fields so you can confirm the additions.

- b. On Available Columns, locate the fields you added to the report query. Check the boxes corresponding to the added columns.

Next, go to ["Add a Server Property for the New Report URI" below](#) to complete your customized case report.

Add a Server Property for the New Report URI

The default URI for case reports is indicated on the Selected Case Report's Parent Groups field:

```
/All Reports/ArcSight System/Core/
```

Your customized case report needs to use a different URI, as described in the following instructions.

To enable your case reports to use the new report format:

Add the full URI, including the customized report name, in `server.properties` through this statement:

```
case.report.uri=<New report URI>
```

For example:

```
case.report.uri=/All Reports/Personal/CaseWorker3's Reports/CW3 Selected Case Report
```

For instructions on setting server properties, refer to the "Managing and Changing Properties File Settings" topic in the [ESM Administrator's Guide](#). After you follow the instructions in that guide, this URI appears on the customized case report's Parent Group field.

Using External Case Management Systems

ESM provides the ability to integrate with ServiceNow®. For other external management systems, you can export ESM cases as XML files.

For information about configuring the integration with ServiceNow®, see the [ESM Installation Guide](#).

To export an ESM case as an XML file:

1. On the **Resources** tab in the **Navigator** pane, select **Cases**.
2. Right-click a case and select **Export > Export to External System**.

The output file is stored in the Manager's archives/exports/ directory. You must configure the external case management system to consume the XML file.

Exporting Cases to ServiceNow®

You can export ESM cases and create them as tickets in ServiceNow®. After you configure integration with ServiceNow®, ESM adds a **ServiceNow® ITSM ID** field to the Console Case Editor. The ID value in this field is a live link to the case. Clicking the ID value launches ServiceNow®.

ServiceNow® manages the session and session expiration between the ArcSight Console and ServiceNow®. If your ITSM session expires and you are not done with your tasks, log in again.

Ensure that your environment meets the following prerequisites before you attempt to export a case:

- Verify that ESM is configured to integrate with ServiceNow®. If not, run the Manager setup wizard and provide the requested information. For more information, see the [ESM Installation Guide](#).
- ServiceNow® must be running and you must have login privileges.

To export an ESM case:

Displaying the case editor is not required for this procedure. You can export a case whether it is locked or unlocked.

1. On the **Resources** tab in the **Navigator** pane, select **Cases**.
2. Right-click the case and select **Export > Export to ServiceNow® ITSM**.
3. Use one of the following methods to log in to ServiceNow®:
 - Use the ServiceNow® Global ID.
 - Provide your ServiceNow® credentials and click **Next**.
 - To use token authentication, select the **OAuth** checkbox and click **Next**. Enter the token and click **Next**.
4. Provide the requested information, and then click **Create Ticket**.

To access the exported case in ServiceNow® :

1. In the ArcSight Console, display the Edit panel for the exported case.
2. In the **ServiceNow® ITSM ID** field, double-click the ID value.
3. Provide the ServiceNow® login credentials.

Chapter 22: Integration Commands

Integration commands leverage the power of security and event management, and broaden its view to show external, snap-in views from ArcSight Logger as well as third-party applications.

ArcSight ESM is shipped with standard content (pre-built commands) and a platform for building your own command configurations.

Contact ArcSight Professional Services if you need assistance in authoring tools integrations with ArcSight products or other applications.

What are Integration Commands?

Integration commands enable you to link from the ArcSight Console to information in other views and applications. You can also build and launch commands locally and on remote servers or appliances, using field values in events as command parameters. You can configure the commands as context-aware, right-click options on different views, resources, and editors on the ArcSight Console.

Configurations can define valid data types and selections for a set of commands. For example, you could configure a set of URL commands to run as a right-click on a selected cell in an active channel and accept only IP addresses as data types.

The ability to integrate commands for various applications means the ArcSight Console can serve as a central hub for defining, managing, and launching Logger searches for older versions of ESM, and third party applications, as well as local ArcSight scripts.



Tip: You can use the ArcSight Command Center to run Logger searches. Refer to the ArcSight Command Center User's Guide, topics on creating peers (which includes ESM and Logger peers) and running searches on those peers.

Supported Command Types

You can build these types of context-based, right-click commands into the ArcSight Console:

Command Type	Output Results	Examples
URL commands link to Web page URLs or URIs	Web browser	Logger searches Tip: Consider using Logger peers for Logger searches using the ArcSight Command Center.
Script commands run scripts	Script/executable output result (for example, action)	Network Tools

Command Type	Output Results	Examples
Connector commands are derived from associated nodes or applications	Structured result based on the SmartConnector and its associated node or application	Refer to your device-specific SmartConnector guide

For more information on working with commands, see ["Defining Commands" on page 668](#).



Tip: All *integration commands* are designed as *manual*, right-click options in various contexts in the ArcSight Console. This enables you to launch commands in ArcSight Console and access available work flows in other applications.

To define *rule-driven commands*, configure rule actions to send *SmartConnector commands* rather than by creating integration commands.

For more information about using SmartConnector commands in rules, see ["Execute Connector Command" on page 541](#).

Local Scripts and Commands to Other Applications

Typical activities for which you might build and run commands in the ArcSight Console that connect to other applications and tools include:

- Launch third-party Web interfaces
- Launch scripts
- Run external searches
- View submitted tickets
- Get Asset/Vulnerability information
- Get Payload Information

You can set up context-aware commands to third-party applications and custom scripts. With command configurations, you can make these available in specified ArcSight Console views and use particular fields as parameters to your commands.

ESM ships with standard utilities configured to be available in ArcSight Console views. For example, the **ping** command is available in grid views such as active channels, lists, and query viewers, and takes as a parameter the *IP address* or *host name* in the selected event.

For information on integrating basic network tools such as Ping, Nslookup, or ArcSight specific Send Logs, see ["Using the Network Tools" on page 73](#) and ["Network Tools as Integration Commands" on page 692](#).

How Integration Commands Work

Integration commands provide resources for tools integration authors to:

- Build context-sensitive commands that you can run locally or on multiple, remote target servers, and you can mix, match, and re-use with configurations.
- Associate parameters with commands to read the resources for which you call the commands. Command parameters make use of Velocity Expressions to pick up values from fields and resources. (See ["Velocity Templates" on page 1088.](#))
- Define configurations sets of commands) for various external applications to specify relevant contexts, commands, rendering formats, and, optionally, remote targets.

Once integration commands and configurations are in place, analysts and operators working with the ArcSight Console can use your custom-built commands or ArcSight pre-built commands for Logger to manage and monitor networks and assets with an extended reach into other views, toolkits, and servers.

Configure Login credentials for authentication on external applications through integration parameters on the user resource. See ["Setting User Login Parameters" on page 683](#) and ["Setting Logins and Other Parameters to Prompt for Values at Runtime" on page 684.](#))

Planning Checklist and Workflow

Plan your command integrations by identifying the utilities or applications to integrate and collecting the necessary information. Here is a checklist of considerations.

Components	Questions
Commands	• What commands will you run on the external application? Is there a subset of commands you want to integrate into the ArcSight Console? • What is the command type (Web URL, local executable script, or Connector command) and syntax?
Servers, Authentication SmartConnectors	Integrating Logger URL commands requires an IP address or Host name of the appliance and authentication credentials for users.


Components	Questions
Configurations	• How do you want to render (display) output results of commands? This largely depends on the command type; for example, URL commands are rendered in an external browser. • How many integration configurations do you need? • Does the application you are integrating have more than one type of interface? (for example, Web and CLI) If so, you'll need a configuration for each interface and associated command type.
Users	• Which users work with these integration tools or applications? • Are authentication parameters required on target servers, appliances, or applications? If so, collect or establish user names and passwords for users who run these commands. • Plan for configuring integration parameters on user accounts for users who work with the external applications. These users need login credentials for both ArcSight and the target applications. • For users with the same authentication parameters for a target server, you can create a target resource with those parameters instead of duplicating the parameters in each user account. Then you can configure the ACL of that target resource so that only those users have access to it. When a command is triggered in the right context, only the target to which that user has access is displayed. Use a similar ACL approach for commands. For example, a single configuration can contain groups of commands, where some commands require special privileges.

Once you have a plan, you might try configuring the commands and testing in this order:

1. Add the commands (command name, type, the command itself, and its parameters).
2. Specify the targets (remote servers where commands run), if any.
3. Create one or more configurations, add in the commands you created, choose how command results are rendered (displayed), and define ArcSight Console UI contexts where these commands are available for use.
4. Add Integration Parameters to User Accounts. If authentication is required on target servers, configure login credentials on user accounts for users who run these commands. These users need login credentials for both ArcSight and on the target applications.
5. Test the commands. See ["Running Integration Commands" on page 685](#)

Navigating to Integration Command Resources

To create or edit integration commands and configurations, start by navigating to **Integration Commands** resources. Following is an example:



Users can access existing integration commands and configurations through right-click commands on the ArcSight Console in various contexts. The contexts depend on how the commands are configured.

Defining Commands

Use the **commands** feature to configure URL, Script, and Connector commands for custom and third party applications and other ArcSight products. Setting up **commands** is the first step in a multi-part process to providing a set of integration commands. (Other tasks include setting up configurations, targets, and user login parameters). This topic explains how to add and edit the command portion of an integration command solution.

Where: Resources > Integration Commands > Commands tab

To add a new command:

1. Right-click a group (folder) in which to create the command, and select **New Command**. This launches the Command Editor in the Inspect/Edit panel. It is best to create new content in your own folder.
2. On the **Command Editor**, select the command **Type**.

Command Type	Description
Script	Executable script that runs <i>locally</i> to the ArcSight Console where the command is launched.
URL	Web URL for which you can define parameters.
Connector	Commands for ArcSight SmartConnectors.

- Required: enter a name.
- Set additional attributes relevant to your command type in the **Integration Target** section of the editor.
 - See ["Script Commands" below](#).
 - See ["URL Commands" on page 671](#).
 - See ["Connector Commands" on page 672](#).

The screenshot shows the 'Inspect/Edit' window for a command named 'Google Search'. The 'Type' is set to 'URL'. The 'Integration Target' section includes fields for Name (Google Search), URL (http://www.google.com/search?q=\$selected), User Attribute, and Password Attribute. The 'Common' section includes Resource ID, External ID, Alias (Display Name), Description, Version ID, and a Deprecated checkbox. The 'Assign' section includes Owner and Notification Groups. The 'Parent Groups' section shows 'admin's Integration Commands'. The 'Creation Information' section includes Created By (admin), Creation Time (15 Nov 2011 09:34:14 PST), and Time Since Creation (1 hour(s) 52 min(s)). The 'Last Update Information' section includes Last Updated By (admin), Last Update Time (15 Nov 2011 09:34:14 PST), and Time Since Last Update (1 hour(s) 52 min(s)). At the bottom, there is a 'Name' field with the prompt 'Enter a name for this command' and buttons for 'OK', 'Cancel', 'Apply', and 'Help'.

- Click **OK**.

Script Commands

ESM users run the ArcSight Console on many different machines. Integration script commands always run on the same machine as the ArcSight Console used to launch them. Therefore, the working directory and program path names should reflect where commands are found in ArcSight Console users' environments

Attribute	Description
Name	User-friendly Name for the command.
Working Directory	Directory containing the executable script. For example, <code>\$systemRoot\system32\</code> You can enter the directory path in the Program field, or click the Browse Directory button  to get a file browser. Use the file browser to navigate to and select the command. Note: Be sure this path reflects the location of the script on machines used by ArcSight Console users for whom you are building these commands.
Program	Full path to the executable command. For example, <code>\$systemRoot\system32\ping.exe</code> You can type the full path to the command in the Program field, or click the Browse Directory button  to get a file browser. Use the file browser to navigate to and select the command. Note: Be sure this path reflects the location of the script on machines used by Console users for whom you are building these commands.
Parameters	Provide parameters for the command. (See "Adding and Editing Command Parameters" on page 673.) The Attributes list provides Velocity Expressions for all event fields and an option to add <code>\$selectedItem</code> as an attribute.



Tip: Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725.](#)

URL Commands

Attribute	Description
Name	User-friendly name for the command.
URL	The URL for the command, along with any parameters provided as arguments to the URL. Click the browse button  to get the Parameters dialog. (See "Adding and Editing Command Parameters" on page 673 for information on how to add the URL along with parameters or <i>arguments</i> to the URL.) You can copy/paste URLs onto the Parameters dialog scratch pad or type them directly. The Attributes link provides Velocity Expressions you can add as parameters (attributes) to the URL. • Type or paste URL directly in the Parameters dialog scratch pad. • Click Attributes to add a Velocity Expression as a URL parameter. Determine the URL by first accessing it from a Web browser address bar. This also shows you where in the URL to add the parameters (if any). Example: Web Search To set up a Google Search on a parameter, do a Google Search in a Web browser. Extract the first part of the URL (everything <i>to the left of</i> the search term) from the Address bar, and paste it into the Parameters dialog scratch pad: <code>http://www.google.com/search?q=</code> Click Attributes on the Parameters dialog to get a list of Velocity Expressions. Select the option, Selections > \$selectedItem. The expression is added as a parameter to the search: <code>http://www.google.com/search?q=\$selectedItem</code> Click OK to close the Parameters dialog and save your changes. Click Apply or OK on the Command Editor when you are satisfied with all settings.) When this search command is deployed as part of an integration configuration, and run via a right-click command in the context of the ArcSight Console, it searches the text in the cell (Viewer table cell) the user selects in the ArcSight Console.



Tip: Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

Connector Commands



Tip: Prerequisites for Connector Commands

If you plan to build and use **Connector commands** you need:

- Access to relevant servers where the SmartConnectors are installed.
- One or more of the associated SmartConnectors deployed and registered with the Manager to which your ArcSight Console is connected.

Test connectivity and authentication between your local machine, SmartConnectors, and appliances before setting up Connector integration commands.

Attribute	Description
Name	User-friendly Name for the command.
Group	Choose a group from the Group drop-down menu. Depending on which Group you select, relevant commands are provided in the next field (Command). Note that CounterACT is no longer supported, and therefore no relevant commands are displayed for it. See Set Event Field in Rule Actions Reference.
Command	Choose a command from the drop-down menu. Depending on which Group you selected, relevant commands are provided here. Choose a Connector command from the drop-down list. Note: In order to get the list of Connector commands, you need to have the SmartConnector deployed and registered with the Manager to which your ArcSight Console is connected.
Parameters	To define parameters for the command: 1. Click the browse button  to get the Parameters dialog. A table of name-value pairs is provided that represents the valid parameters for the given command. 2. Select the parameters to use, and define values for them with either hard-coded values or <i>Velocity Expressions</i>. For example, you could define the Connector command Quarantine Node By IP Address to use three parameters; IP Address, Quarantine Period, and Overwrite Active Quarantine (a yes/no value set to 0 or 1, respectively). You could set the IP address to a Velocity Expression for <i>attacker address</i>, Quarantine Period could be set to 1 hour, and overwrite set to Yes. The Attributes list provides Velocity Expressions for all event fields along with options to add Console selections, dates, and channel start and end times as attributes. 3. Click OK on the Parameters dialog to save your changes.



Tip: Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see "[Common Resource Attribute Fields](#)" on page 725.

Adding and Editing Command Parameters


The Attributes list includes Velocity Expressions for all event fields and an option to add user field or item selections, channel start or end time, date/time, and other Velocity Expressions as attributes.

Where:

- For Script of Connector commands, in the Parameters field
- For URL command, in the URL field


To provide parameters for a command:

1. Click the browse button  to get the Parameters dialog.
2. Click **Attributes** to get a list of variables and Velocity Expressions. The right-click menus on this panel give you access to velocity expressions, fields, and variables you can use as parameter values in commands.



3. Select the expression you want to add. The attribute list includes Global Variables. If the global variable to add is composed of a list of fields, expand the global variable displayed in the Parameters dialog and select the field you want.

The expression is added to the Edit Attributes scratch pad as a parameter.



4. You can continue adding expressions, which are chained together.

For example, selecting **Threat > Priority** from the Attributes list results in this parameter being placed on the scratch pad:

```
${priority}
```

Subsequently selecting **Attacker > Address**, updates the scratch pad entry with chained-together expression:

```
${priority} ${attackerAddress}
```



Tip: The Parameters dialog is an editable note pad. In addition to adding Velocity Expressions from the Attributes menu and Templates for Connector command parameters, you can type new expressions directly into the dialog. Also, you can select and edit existing expressions manually.

See also ["Removing a Command Parameter" below](#).

5. When the Parameters scratch pad reflects the expressions you want to include as command parameters, click **OK**.

The parameters you added are reflected on the Attributes tab in the Command Editor.

Parameters \$ (priority)

Be sure to click **Apply** or **OK** on the **Command Editor** to save changes to command parameters along with any other changes to the command that you want to retain.

Removing a Command Parameter

Where:

- For Script of Connector commands, in the Parameters field
- For URL command, in the URL field

To remove a command parameter:

1. Click the browse button  to get the Parameters dialog.
2. Select the parameter in the scratch pad and hit the Delete key on your keyboard.
3. To add a new parameter to replace the one you are deleting, do so by following steps described in ["Adding and Editing Command Parameters" on page 673](#).
4. Click **OK** on the Parameters dialog.
5. Click **Apply** or **OK** on the **Command Editor** to save your changes.

Using Configurations to Group Commands

An integration **configuration** resource represents a family of commands of the same type. Commands in a configuration share the same context, rendering method, and targets.

Configurations provide a way of grouping similar commands and specifying common options for where on the ArcSight Console UI the commands are available (**contexts**) and where commands run (scripts run locally; others, like Connector commands, can have one or more remote **targets**). This is partly a matter of preference (about how you want to group, organize,

and present commands to ArcSight Console users), and partly a matter of which commands belong together.



Note: *Configurations can include only commands of the same type* (script, URL, or Connector). Commands that share a configuration use the same renderer, contexts, and (if relevant) targets. You might want to make finer-grained groupings; for example, sub-groups of scripts or Connector commands.

Setting up *configurations* is a step in a multi-part process of making a set of integration commands available to ArcSight Console users. (Other tasks include setting up commands, targets, and user login parameters).

This topic explains how to add and edit the *configuration* portion of an integration command solution. For an overview of the integration commands feature, see ["Integration Commands" on page 664](#). For more details on the relationship between commands, configurations, and targets, see ["How Integration Commands Work" on page 665](#).

To create a configuration:

1. In the Navigator panel, select the **Integration Commands** resource from the drop-down menu and click the **Configurations** tab.
2. Right-click a group (folder) where you want to create the configuration, and select **New Configuration**.

This launches the Configurations Editor in the Inspect/Edit panel.

3. Select a type for the configuration you are defining:
 - **Script** (see [Script Commands](#))
 - **URL** (see [URL Commands](#))
 - **Connector** (see [Connector Commands](#))
4. Fill in the fields on Attributes, Context, Commands, and Targets tabs as described in:
 - ["Configurations Attributes" on the next page](#)
 - ["Configurations Contexts" on page 678](#)
 - ["Configurations Commands" on page 679](#)
 - ["Configuration Targets" on page 680](#) (when commands run on remote targets)
5. Click **Apply** or **OK** to add the new configuration.



Tip: You can use the above procedures to create integration configurations from within a context. To do this, right-click anywhere in the UI, and choose Integration Commands > New Configuration. Then perform the above steps.

Configurations Attributes

Define the configuration name and other basic details for the configuration on the Configurations **Attributes** tab.

Configuration Attributes	
Attribute	Description
Type	Choose the type of configuration from the drop-down menu: • Script • URL • Connector Note: The configuration type must match the command types in the configuration. (See "Defining Commands" on page 668 for configuration instructions.) Once the configuration is saved, the type is not editable. This setting influences choices on other options for the configuration, such as the "Type" above.
Name	A user-friendly, informative name for the configuration that (preferably, one that indicates the commands contained in it).
Allow Multi Select	Use this to allow selecting multiple events on which to run a command. It is off by default. A check mark indicates it is on/enabled. When on, users can select multiple events and the commands assign the values to a parameter as a comma-separated list. For example, suppose you have a command with the parameter <code>ip=\$targetAddress</code>. • With Multi Select disabled, the command accepts only a single IP address based on a selected event (for example, <code>ip=192.0.2.0</code>). • With Multi Select enabled, a user can also get <code>"ip=192.0.2.1,192.0.2.2"</code> if two rows are selected. In order for this to work: (1) the ArcSight Console context (for example, active channel) must allow multi-row selection, and (2) the integration target must support a comma-separated list of values for the given command and parameter. Note: Multi Select does not affect how individual fields in an event are processed. Event field processing is determined entirely by the definition of command parameters. For example, a command with an Attacker Address parameter always gets that value from the selected event.



Tip: Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

Configurations Contexts

As a part of constructing command configurations, you can configure **contexts** for where in the ArcSight Console certain commands are available. At the same time, you can define **parameters** for picking up and passing the value in any selected cell, row, or event field.

For example, you could configure a URL command for a Google search as a right-click command on any cell in an ArcSight Console grid view. By using a parameter as the argument to the search command, you could pick up the text from the selected cell or value from any selected field to use as your search term. (In the Commands editor, all fields, provided as a list of Velocity Expressions, are available for use as command parameters.)

Once configured, integration commands are available on right-click context menus from a variety of contexts including:

- Relevant fields in active channels (for example, IP address, host name, MAC address)
- Relevant resources (for example, assets)
- Active Lists, sessions lists, query viewers and channels

Also, you can configure **user login parameters** on ArcSight Console users (via a new Integration Parameters tab in the Users resource editor), thereby binding user login information to commands for third-party or ArcSight applications that require secure logins. (See ["Setting User Login Parameters" on page 683](#) for more information.)

You can configure a command to prompt for parameter information, which is often useful for login scenarios and as well as others. (See ["Setting Logins and Other Parameters to Prompt for Values at Runtime" on page 684](#) for more information.)

To set up command contexts:

Use controls on the Configurations **Context** tab to add, edit, or remove contexts in a configuration.

Configurations

AttributesContextCommandsTargetsNotes

AddRemove

Location	Type	Selection	Data Type
Editor	All Editors	All Selections	All Data Types
Viewer	All Views	All Selections	All Data Types

Click the fields under Location, Type, Select, and Data Type to get drop-down menus with which to select contexts in the ArcSight Console UI where the command is available and to which selections it applies.

Command Context Attributes	
Attribute	Description
Location	View where in the ArcSight Console the command is available. For example: • Viewer, for the Viewer panel where Views of active channels, dashboards, and so on are shown • Resource, for the Navigator Panel resource tree • Editor, for resource editors
Type	Contexts in the ArcSight Console panels where the command is available. Available types vary depending on the location you choose. For example, if you choose Viewer for the location, you can specify types of “views” where you want the command to display, such as Grid View, Chart View, various List entries, Dashboards, Query Viewers, and so on.
Selection	User selection or subset of it that is fed into the command. Options can include All Selections, Selected Cell, Selected Row, Selected Attribute.
Data Type	Data type for the parameters fed into the command (derived from the Selection). Options include: • All Data Types • IP Address • MAC Address • Date • Double • Integer • Long • Resource • String

Configurations Commands

Where: Resources > Integration Commands > Configurations > Commands tab


To add a command to a configuration:

Use controls on the Configurations **Commands** tab to add, edit, or remove commands in a configuration.

1. Click **Add** to bring up the Commands Selector dialog.
2. Navigate to and click (checkmark) the commands you want to add, and click **OK**.

The commands are added to the list. (You can add multiple commands to a single

configuration.)



To edit commands in a configuration:

Select the command you want to edit and click **Edit**.

This provides a shortcut into the **Command Editor** for the selected command. See ["Defining Commands" on page 668](#) for information on editing the command.

To remove commands from a configuration:

On the Configurations **Commands** tab, select a command in the list and click **Remove**.

Configuration Targets

Targets are not required for all command types, only for those that run on remote servers. Before you can add a target to a Configuration (explained here), you first need to define it as described in ["Specifying Targets" on the next page](#).

Use controls on the Configurations **Targets** tab to add, edit, or remove targets in a configuration.



Note: If you plan to add remote targets to a configuration, you need host information for the remote servers and login credentials if authentication is required.

Adding a Target to a Configuration

Targets are applicable to any commands that you want to send to a remote server.

- Click **Add** to bring up the Connectors Selector dialog.
- Navigate to and click (check mark) the target you want to add, and click **OK**.

Editing Targets in a Configuration

- Select the target you want to edit and click **Edit**.
- This provides a shortcut into the **SmartConnector Configuration** Editor for the selected connector or target. (See ["Configuring the SmartConnector" on page 175](#), and also the *SmartConnector User's Guide*.)

Removing Commands from a Configuration

On the Configurations **Contexts** tab, select a target in the list and click **Remove**.

Specifying Targets

Optionally, you can specify targets (remote servers where one or more commands run).

If you have multiple remote servers, you might want to configure multiple targets on which to run a single command with the same or different parameters.

For example, you can configure any of the following as command targets.

- Applications with Web interfaces/clients like ArcSight Logger appliances
 - Search providers (for example, Google, Yahoo, ask.com)
 - IT/Security portals
 - Asset/Vulnerability information
 - Ticketing Web servers
- Connectors

Setting up *targets* is a step in a multi-part process of making a set of integration commands available to ArcSight Console users. (Other tasks include setting up commands, configurations, and user login parameters).

This topic explains how to add and edit the *configuration* portion of an integration command solution. For an overview of the integration commands feature, see ["Integration Commands " on page 664](#).

To add a new target:

1. In the Navigator panel, select the **Integration Commands** resource from the drop-down menu and click the **Targets** tab.
2. Right-click a group (folder) where you want to create the target, and select **New Target**. This launches the Command Editor in the Inspect/Edit panel.
3. Fill in the fields as described below.
4. Click **Apply** or **OK** to add the new target.

Target Attribute

The only target attribute you need to provide is a user-friendly name for the server.

Attribute	Description
Name	Name for the remote server or appliance where the command run.

Target Integration Parameters

Targets are used only for URL configurations, where you parameterize the Web host target of the URL, and sometimes login credentials. Type directly into the fields to define a parameter, as described below.

Field	Description
Parameter	Parameter name, as specified in the command definition related to this target.
Type	Parameter type. Choose Text or Password from the drop-down menu. Password type parameters are automatically encrypted. Notes: • Always set login credentials (passwords or authentication tokens) to type “Password” (not “Text”). (Credentials set to “Text” are not masked on the UI and are sent as clear text if the renderer is an external browser.) • You can set passwords and authentication credentials on target servers too, but we recommend against it in most cases. Doing so risks opening up a target server to any user who has access to the integration commands (not necessarily an account on the target server). Additionally, it does not give you any tracking information based on user logins to the server.
Value	Hard-coded value, variable, or Velocity Expression for the parameter. For example: • A host name or IP address as a value for a target server parameter

To add a new parameter, click **Add**. This gives you a new row in which to enter Parameter, Type, and Value information. You can add multiple parameters to a target.



Tip: Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

Authorization and Authentication Settings

Authentication: You can specify user login behavior for commands designed to run on secure, remote target servers. You can specify login credentials to be used as part of the command, or set parameters that prompt users to enter their user name and password when they run the command.


Authorization: You can set up fine-grained access control lists (ACLs) to specify which ArcSight Console users have permissions to view, run or edit different commands. See ["Managing Permissions" on page 107](#).

Setting User Login Parameters

You can specify login credentials on user accounts or on remote target servers. It is best to set login credentials on user accounts, but both options are described below. (Login credentials are not required for Connector integration commands because the authentication is handled as part of the SmartConnector setup.)

Setting Login Credentials

For URL commands on remote targets and script commands that run locally, you can define login credentials as a part of user configurations. (Choose Navigator > **Users**, select and edit a user or create a new one, then click the **Integration Parameters** tab on the **User Editor**.)



Defining login information as part of user accounts gives you the flexibility to configure multiple users, each with different logins. In this case, login credentials are not tied to the command target, but rather associated with individual users.

A single user account can have login credentials for different servers and scripts. .



Tip: For **security best practices**, we recommend that you:

- Always set login credentials (passwords or authentication tokens) to type "Password" (not "Text"). Credentials set to "Text" are not masked on the UI and are sent as clear text on the Web browser.
- Save authentication information only as parameters on user accounts, not on target servers. This strategy binds authentication details to specific users, and gives you tracking information based on user logins (for example, you can tell which users ran which commands and when).

Examples of authentication information are user name and password combinations, and authentication tokens sent in URLs.

Setting Login Credentials on Target Servers

Although not generally recommended, login credentials for URL commands on remote targets also can be defined as part of the Target definition, as described in ["Specifying Targets" on page 681](#). (Choose Navigator > Integration Commands > Targets tab, select and edit a target or create a new one, then click the **Integration Parameters** tab on the **Targets Editor**.)

If login information is defined here, everyone who uses the command uses the same credentials to log in to the remote target server.



Caution: Here are best practices to secure your authentication information.

- Do not save authentication information as parameters on target servers. It runs the risk of opening up a remote server to any user who has access to the integration commands. Additionally, it does not give you any tracking information based on user logins to the server.
- Always set login credentials (passwords or authentication tokens) to type **Password**, not Text. Credentials set to Text are not masked on the UI and are sent as clear text if the renderer is an external browser.)

Setting Logins and Other Parameters to Prompt for Values at Runtime

You can set parameters for which you would like to prompt users to specify values at runtime, such as user name and password, host names, IP addresses, and other options.

When an integration command runs (that is, when a user selects an integration command in some context on the ArcSight Console), the command first looks for any required parameter values in a variety of sources, including in the command statement itself, in the defined context, on the user account, on the target (if there is one), and so forth. If it doesn't find parameter values in any of these places, the system prompts the user to type in the values.

You can include login and other parameters as flags on a script command that runs against a server, as shown here for the archive command which runs on a Manager. When this command is run, it prompts the user for a Manager host name and administrator password. (It does not prompt for the user name, admin, since this is already provided in the command statement.)

```
archive -action import -m $hostmgr -u admin -p $passwd -f abc.arb
```

Refer to ["Entering/Saving Command Parameters at Runtime" on the next page](#) (in ["Running Integration Commands" on the next page](#)), for an example of the run-time prompts users see when they run this command.

Running Integration Commands

After commands are configured, they are available in various contexts in the ArcSight Console.

For example, suppose you have a configuration for a set of commands with the contexts set as follows:

Location	Type	Selection	Data Type
Viewer	All Views	All Selections	IP Address

This means that the given commands are available on right-click context menus on any view (for example, active channels, list views, chart views, dashboards, and so on). The user can select any row, cell, or area on a chart. In this context, only IP addresses can be provided as valid parameters to the command.

Open an active channel, session list, active list, dashboard, or other resource in the viewer that shows, for example, a suspicious device, machine, or user that you want to quarantine.

Find the row on the Viewer display that contains the suspicious entity, and select a cell in that row that contains the source IP address (for example, Attacker Address).


Right-click over the cell with the source IP address (for example, Attacker Address), and choose **Integration Commands > Quarantine Node**.

This launches the selected command, using the IP address for the selected cell as the parameter for the command.

In general, a right-click any context in the ArcSight Console UI for which integration commands have been configured show all integration configurations.

Entering/Saving Command Parameters at Runtime

Commands can be configured to prompt for parameter values at runtime (as described in ["Setting Logins and Other Parameters to Prompt for Values at Runtime" on the previous page](#)). Also, if ready-made commands (such as for Logger) are not pre-configured, you are prompted for values. For example, parameters might ask for a particular host name as command input, an IP address against which to run a command, or login credentials to a target server.



The image shows a 'Parameters' dialog box with a table of parameters. The table has five columns: 'Save To Target', 'Save To User', 'Parameter', 'Type', and 'Value'. The parameters listed are OTPPassword, LoggerHost, OTPUser, LoggerPort, LoggerUser, and LoggerPassword. The 'LoggerPassword' row is highlighted, and its 'Type' is set to 'Password' with a dropdown arrow. The 'Value' field for 'LoggerPassword' contains '*****'. The 'LoggerUser' row has 'admin' in the 'Value' field. At the bottom of the dialog are 'OK', 'Cancel', and 'Help' buttons.

Save To Target	Save To User	Parameter	Type	Value
☐	☐	OTPPassword	Text	
☐	☐	LoggerHost	Text	
☐	☐	OTPUser	Text	
☐	☐	LoggerPort	Text	
☐	☒	LoggerUser	Text	admin
☐	☒	LoggerPassword	Password	*****

If you launch a command that prompts for input, enter the appropriate text in the “Value” field for each required parameter.

If you have appropriate permissions, you have the option to save parameter values with the target or with your user account so that you don’t have to re-type them each time you run the command.



Tip: Security Best Practice Recommendations:

- In order to save parameter values at runtime, you need to belong to a group with *read and write permissions* to the associated *targets*.
- Always set login credentials (passwords or authentication tokens) to type “Password” (not “Text”). Credentials set to “Text” are not masked on the UI and are sent as clear text if the renderer is an external browser.
- Always save login credentials to user (Save to User), not to the target server. This strategy binds authentication details to specific users. This better safeguards access to the remote server to appropriate users. Also, you have tracking information based on logins as to which users entered which commands.

If you save authentication details as parameters on a target, you run the risk of opening up a remote server to any user who has access to the integration commands (but not necessarily an account on the target server). And you have no per-user tracking information.

Using the Recon Integration Commands

Integration with Recon requires specific browser versions. See the *ESM Support Matrix* for details.

If a Recon deployment is available, you can create target parameters then execute commands to run a search. You can run searches:

- By Source and Destination
- By Vendor and Product

Where: Navigator > Resources > Integration Commands

Configure target with target parameters:

1. Select the **Targets** tab.
2. Expand /All Integration Targets/ArcSight Administration/Recon.
3. Right-click **Recon 1** and select **Edit Target**.
The target's Edit panel opens.
4. Click the **Integration Parameters** tab.
5. Click the **Add** icon and add the first parameter using the following:
Parameter name = **Host**
Type = Text (the default)
Value = Enter the host information for Recon
6. Click the **Add** icon again and add the second parameter using the following:
Parameter name = **Port**
Type = Text (the default)
Port = **443**, the default port for Recon. If using a different port, enter it here.
7. Create additional targets if applicable, for example, using a different value for host or port.

Run the search commands:

1. Open an active channel.
2. Right-click an event, select **Integration Commands**, and select **Recon Search**.
3. In the popup, select a type of Recon Search command:

By Source and Destination

or

By Vendor and Product

Select a target, then click **OK** to close the popup.



Note: If you did not perform the previous set of steps in [Configure target with target parameters](#), you are next prompted in another popup to enter the IP address for the Recon host. In the same popup, you have the option to save this IP address parameter to the target. See ["Entering/Saving Command Parameters at Runtime" on page 685](#) for information.

4. Click **OK** again.

Your preferred browser launches the Recon Search page.



Note: On the Recon page, the time range for the search is the last 30 minutes by default, which may not yield any search results. If necessary, edit the active channel by changing the **Start Time** and **End Time** values for your search.

See ["Creating or Editing an Active Channel" on page 224](#) for details on setting those values.

See also ["Running Recon Searches" on page 293](#).

ArcSight Logger Search Commands



Note:

The Command Center provides an easier way to search events in ESM and ArcSight Logger than integration commands.

Refer to the Searches topic in the *ArcSightCommand Center User's Guide* for more information on setting up Logger peers and performing searches on those peers.

ArcSight Logger Search commands are integrated into the ArcSight Console and provided as standard content. These commands enable searches on ArcSight Logger appliances and are supported starting with ArcSight Logger v4.0.

Configuring the ready-made integrated Logger commands involves specifying IP addresses for the target Logger appliances and saving login credentials on users who need Logger search access.

Logger Integration Commands



Note:

The Command Center provides an easier way to search events in ESM and ArcSight Logger than integration commands.

Refer to the Searches topic in the *ArcSightCommand Center User's Guide* for more information on setting up Logger peers and performing searches on those peers.

These integrated commands are supported in ArcSight Logger v4.0 and later versions. (These are defined in /All Integration Commands/ArcSight Administration/Logger.)

Logger Command	Description
Logger Search	Allows you to right-click an event in an active channel and run a search for one of the fields presented in a list. • Displays a pop-up dialog with search options. • Allows you to search by: - Event Name - Destination - Source - Destination and Source - User - Service Vendor and Product • You can select the Logger appliance on which to search.
Logger Quick Search	Allows you to right-click a field in an active channel to search based on the field and value selected. If there is more than one Logger appliance set up, a pop-up dialog box allows users to choose which appliance to search.

Enabling Integrated Logger Searches



Note:

The Command Center provides an easier way to search events in ESM and ArcSight Logger than integration commands.

Refer to the Searches topic in the *ArcSightCommand Center User's Guide* for more information on setting up Logger peers and performing searches on those peers.

You can use either of these approaches for setting up the Logger searches:

- Configure target and authentication details before the commands are run (for example, a single administrator specifies Logger targets, users, and command parameter values). For this workflow, follow the steps below to:
 1. [Set Up Logger Command Targets](#),
 2. [Set Up the Logger Command Configuration](#), and
 3. [Set Up Users for Logger Access](#).
- Configure commands at command runtime (for example, if you launch a command based on an active channel selection and fill in the target appliance IP address, authentication token, and parameter values on-the-fly).

If you have not pre-configured the command details, you are prompted to enter the right values when you run the commands. Information on how to run integration commands, see ["Running Integration Commands" on page 685](#). (This topic includes information about running commands and entering/saving parameter values at command runtime.)

To enable pre-configured ArcSight Logger searches from the ArcSight Console, follow these steps.

1. Set Up Logger Command Targets

1. In the Navigator, click the **Resources** tab, and then navigate to **Integration Commands > Targets**.
2. Create an integration target for your Logger Appliance, or edit an existing target in /All Integration Targets/ArcSight Administration/Logger/, such as Logger Appliance 1 or Logger Appliance 2.
3. In the Target editor on the **Integration Parameters** tab, add these parameters.

Parameter	LoggerHost
Type	Text
Value	<IP address or Host Name of the Logger Appliance>

4. If you have more than one Logger appliance, create an additional integration target for each appliance to be made searchable.
5. Click **Apply** or **OK** to save your changes to the target.
6. For information about command targets, see ["Specifying Targets" on page 681](#).

2. Set Up the Logger Command Configuration

1. In the Navigator, click the **Resources** tab, and then navigate to **Integration Commands > Configurations**.

You can find the pre-built Logger command configurations in /All Integration Configurations/ArcSight Administration/Logger/

2. Edit the **Logger Search** integration configuration:

In the Configuration editor, click the **Targets** tab and add the integration targets you created earlier.

3. Edit the **Logger Quick Search** integration configuration:

In the Configuration editor, click the **Targets** tab and then add one integration target from the list of targets you just created.

4. Click **Apply** or **OK** to save your changes to the configuration.

For information about command configurations, see ["Using Configurations to Group Commands" on page 675](#).

3. Set Up Users for Logger Access

1. In the Navigator, click the **Resources** tab, and then navigate to **Users**.
2. Edit the users that have access to the Logger appliance. In most cases, these users should have administrator privileges.
3. Click the **Integration Parameters** tab, and then specify the integration parameters for the Logger search

Parameter	Description
LoggerUser	The Logger account name to use when accessing a Logger target.
LoggerPassword	The password for that Logger account.
LoggerHost	The IP address of the Logger host.
OTPUser	The Logger user account to use with the OTP authentication feature (see note, below).
OTPPassword	The password to use for the OTPUser specified above.
LoggerPort	For OTP, you must specify the port number for the Logger. For a Logger appliance the port number is 443.



Note: The **One-Time Password (OTP)** function requires that you are using Logger 5.1 or later.

If you have an earlier version of Logger, then for Logger searches continue to use the LoggerUser and LoggerPassword as before. Until you install Logger 5.1 or later, searches display a message that it failed to negotiate a single-use session token and is proceeding with regular authentication. Just click **OK** to continue.

For all of these parameters, you can specify a type of “Password” or “Text.” For Text type, the password appears in plain text in the session URL. This is not an issue for any parameter except passwords, so use text for the others but use the Password type for passwords.

Save to Target means that this parameter automatically applies whenever any user is accessing the currently-specified target Logger.

Save to User means this parameter is used automatically whenever the currently-logged-in user accesses any Logger target.

4. Click **Apply** or **OK** to save your changes to the user.

For information about setting up login credentials and access control lists for integration commands, see ["Authorization and Authentication Settings" on page 682](#).

Example of Running a Logger Quick Search



Note:

The Command Center provides an easier way to search events in ESM and ArcSight Logger than integration commands.

Refer to the Searches topic in the *ArcSightCommand Center User's Guide* for more information on setting up Logger peers and performing searches on those peers.

With Logger target appliance information and user authentication details, you are ready to run Logger searches from the ArcSight Console.

Right-click on an event in an active channel, select a field such as Device Vendor or Device Product, and choose **Integration Commands > Logger Quick Search**.

This logs you into the target Logger appliance (for example, Logger Appliance 1), and sends values for the selected fields as parameters on the search command.

You get the results of the search in the browser, and are now logged in to the Logger Web client.

Network Tools as Integration Commands

The following standard network tools are also provided as integration commands. You can find this toolset in /Integration Commands/Shared/ArcSight System/Tools/. You can edit these or add new commands, configurations, and contexts as described in ["Defining Commands" on page 668](#) and ["Using Configurations to Group Commands" on page 675](#). (Also see ["Using the Network Tools" on page 73](#))

With network tools integration commands you can:

- **Define contexts for where tools show up on the ArcSight Console.** You can customize integrated network tools and configure them for all types of views (charts, graphs, tables), and in the navigator, editors, and so on. Legacy network tools are available only on grid views; you cannot define the context.
- **Select and run commands on navigator tree items, all types of views, and editors items.** With integrated network tools, you can select various items in chart and graph views, on the editors, and in the navigator tree. Legacy network tools are limited to running only on the selected cell in a grid view (table) in the Viewer.
- **Configure access control lists (ACLs).** You can grant or limit access to integrated network tools commands for particular user groups by setting the setting ACL permissions on the tools resource group. The integrated network tools reside under /All Integration

Configurations/ArcSight System/Tools. Under Tools, the network tools are further grouped into Linux and Windows.

You can control access to the tools commands and configurations groups (select the Tools group, right-click, and choose **Edit Access Control**) as described in ["Granting or Removing Resource Permissions" on page 108](#). You can organize users and the tools themselves into various groups to fit with the permissions scheme you want to create.

Tree	Resource
Nslookup	Resolves an IP address to a host or domain name, or the reverse. Notes: - For the Linux version of this command, Nslookup(Linux) is for IPv4 address or hostname, and Nslookup-IPv6 (Linux) is for IPv6 address or hostname. - For the Windows version of this command, Nslookup (Windows) works for both IPv4 and IPv6 addresses or hostnames.
Ping	Determines whether a particular IP address is online and/or it tests and debugs a network by sending a packet and waiting for a response. Notes: - For the Linux versions of this command, Ping (Linux) is for use in an IPv4 network, and Ping6 (Linux) is for use in an IPv6 network. - For the Windows version of this command, Ping (Windows) works for both IPv4 and IPv6 networks.
PortInfo	Lists standard usage such as WWW or FTP, for a specified port number. Note: This command works in both IPV4 and IPv6 networks.
Traceroute	Shows the path from the ArcSight Console to the IP address selected in the grid view, reporting the IP addresses of all routers in between. Note: This command works on both IPV4 and IPv6 addresses.
WebSearch	Search the Web through Google to find links to the keywords present in currently selected active channel grid view cells. Note: This command works in both IPV4 and IPv6 networks.
Whois	Looks up who is behind a given domain name; information might include addresses and telephone numbers. Note: This command works in both IPV4 and IPv6 networks.

These are configured with default Velocity Expressions for parameters. You can edit the commands and configurations for these network tools as needed (and add new ones).

To run a network tool, select an IP address in a grid view (for example, active channel, list, data monitor) and select **Integration Commands > <Network Tool>** from the context menu (for example, **Integration Commands > ping**).



Note: The Send Logs command is not configured as an integrated command. See ["Using the Network Tools" on page 73](#) and ["Send Logs" on page 1047](#) for information on that command.

To add or reconfigure legacy tools:

1. Choose **Tools > Local Commands > Configure**.
2. Select a tool and click **Edit**.

Keep in mind that they have limitations compared to the new tools.

More Integration Examples

To experiment with building integration commands, you need one **command** and one **configuration**. Create the commands first because the configuration references the commands.

The configuration also defines how command results are rendered, and references **contexts** where your new Integration Commands appear in the ArcSight Console right-click menus (for example, Viewers, Resource Panel, Editors, and more specifics within those contexts).


To define **targets** (remote servers where commands run), add them to the configuration.

Here are examples of how to set up a command to do a Google Search on a selected cell in the ArcSight Console, and how to set up commands that use Google Maps to locate a target and an attacker. The examples do not require a “target,” so just set up a command, add it to a configuration, and run it. The details of this and other types of commands and configurations are discussed further in the topics that follow.


To add a command for Google Search:

1. Start by getting the format of the Google search. Do a Google Search in a Web browser. Copy the first part of the URL (everything *before* or *to the left of* the search term) from the Address bar, so you have it on your clipboard. (You paste in to the Parameters dialog in a later step.)
2. In the ArcSight Console Navigator panel, select the **Integration Commands** resource from the drop-down menu and click the **Commands** tab.
3. Right-click the group (folder) where you want to create the command and select **New Command**.
4. On the Commands Editor, fill in these attributes:
 - For command Type, choose **URL**.
 - For Name, provide a user-friendly name like **Google Search**.

- For URL, click the browse button  to get the Parameters dialog. Paste the Google search prefix into the Parameters dialog scratch pad:
`http://www.google.com/search?q=`
- Click **Attributes** on the Parameters dialog to get a list of Velocity Expressions. Select the option, **Selections >\$selectedItem**. The expression is added as a parameter to the search: `http://www.google.com/search?q=$selectedItem`




- Click **OK** to close the Parameters dialog and save your changes.
- Click **Apply** or **OK** on the Commands editor to save the command.




1. Set up the configuration and add the command to it:
 - a. Click the **Configurations** tab.
 - b. Right-click a group and select **New Configuration**.
 - c. On the Configurations Editor, select **URL** as the configuration Type, and enter set these attributes:
 - For Name, provide a user-friendly name.
 - The output will be rendered on a preferred Web browser specified during Console installation.
 - d. Click the **Context** tab. This sets where in the ArcSight Console the command is available. Click **Add** to get a set of context fields, then click into each field to select a location, type, selection, and data type. (You can add multiple contexts by clicking Add again.) Add one context to show in the Viewer in all "views" and to take the selected cell as the "selection":

Location	Type	Selection	Data Type
Viewer	All Views	Selected Cell	All Data Types



When the search command is deployed as part of this configuration, and run using a right-click command in the context of the ArcSight Console, it searches on the text in the “cell” (Viewer table cell) the user selects in the ArcSight Console.

- e. Add the command to the configuration. On the Configuration Editor, click **Commands**. Click **Add** to get the command selector, select your Google Search command, and click **OK**.



- f. Click **Apply** or **OK** on the Configurations Editor to save the configuration.
2. Run the Search command you just built:
 - a. Open any active channel, list, data monitor, or query viewer with a table style view.
 - b. Right-click any cell in the Viewer that contains a term you would like to search on, and select **Integration Commands > Google Search** (or whatever you named the command).

The command runs a search using the text from the selected cell as the search term, and returns search hits in your preferred browser.

To add a command for Google Maps:

Use the same basic steps in the previous example for Google search to integrate a command named **Google Maps**. This Google Maps example lets you pass the GeoLatitude and

GeoLongitude to locate, for example, an attacker and a target. Refer to the following information as a guide for the URLs in your integration command.

Command	URL and attribute
Attacker	http://maps.google.com/maps?q=\${attackerGeoLatitude},\${attackerGeoLongitude}
Target	http://maps.google.com/maps?q=\${targetGeoLatitude},\${targetGeoLongitude}

Chapter 23: Knowledge Base Authoring

These topics explain how to do the basic tasks of managing Knowledge Base articles.

Managing Knowledge Base Articles

You can apply the following methods to incorporate knowledge base (KB) articles in the Console.

- Write a new article in the Console using your preferred text or HTML editor. For this method, make sure you have set your Console preferences for your preferred editor. See ["Setting Default Editors and Viewers" on page 49](#) for the procedures.
- Import an existing file or web page.
- Reference an existing file or web page.

Writing or importing the article means the content is stored in the Console and you edit the content in the Console. If you edit an imported article in the Console, the original remains untouched. Conversely, editing the original of an imported article does not affect the content that resides in the Console. Referencing an article means the content is being maintained externally.

Where: Navigator > Resources > Knowledge Base

To enter or change basic information about the KB:

1. If you are creating a KB, right-click a group and select **New Article**.

If you are editing a KB, right-click the KB and select **Edit Article**.



Tip: In edit mode, the Launch Editor and Upload File buttons are enabled.

2. In the Knowledge Base Editor, select the **Article** tab.
3. On the Article tab, type in the **Name** text field for the new KB, or make edits for the existing KB.
4. Optionally use the **Summary** field to add a brief description of the article.
5. Optionally enter a different name for the information source in the **Author/Credits** field.
At this time, the Upload File and Launch Editor buttons are still disabled for new articles. Note that this condition applies only if you are defining a KB for the first time.
6. In the Import/Reference field, select **Import** to enable the buttons.

From here, you can use the Console to write the article using your preferred editor. You can also choose to import an article or to reference an article.

To write the article using your preferred editor:

This assumes you have completed the steps to enter basic information about the KB.

1. Click **Launch Editor** to start your editor application.
2. In the editor application, type the content of your article. Save the file.
The file has an assigned filename using a default path where the Console is installed. On the Console's edit panel for the KB, this path is inserted in the Origin URL field.
3. Click **Apply** to save the KB definition.
After saving, the Origin URL field is cleared.
4. Verify that the content is added to the KB by clicking the edit panel's **Editor** tab.
If the content needs to be changed in the future, do it here.
5. Click **Preview** if you want to see how the article will appear in your default Web browser.

To render text as HTML:

This assumes you have completed the steps to enter basic information about the KB and that you are familiar with authoring in HTML.

1. Create a KB article or edit an existing article.
2. In the Import/Reference field, select **Import**.
3. In the **Editor** tab, enter your HTML-formatted article using the following general guideline:

```
<html>  
<body>  
Your article here  
</body>  
</html>
```
4. Click **Apply** to save the KB definition.
5. Click **Preview** if you want to test.
6. Read the information in the resulting popup and click **Yes** if you want to continue rendering the article in HTML.

The preview is displayed in your preferred Web browser.

To import an article:

This assumes you have completed the steps to enter basic information about the KB.



Note: When you select **Import**, only the content between the body tags of the HTML page appear. Therefore, you should use **Reference** when the HTML page uses JavaScript, uses frames, or includes images. Select **Import** after you have used an editor to create a file specified in the **Origin URL** text field.

1. In the **Origin URL** text field, type the URL or directory path to the page or file.
2. In the Import/Reference field, verify that **Import** is selected.
3. Click **Apply** to save the KB definition.
4. Verify that the content is added to the KB by clicking the edit panel's **Editor** tab.
If the content needs to be changed in the future, do it here.
5. Click **Preview** if you want to see how the article will appear in your default Web browser.

To reference an article:

This assumes you have completed the steps to enter basic information about the KB.

1. In the **Origin URL** text field, type the URL to the page or file that contains the article's information.
2. In the Import/Reference field, select **Reference**.
3. Click **Apply** to save the KB definition.
4. Click **Preview** if you want to see how the article will appear in your default Web browser.

To show an article:

In the Knowledge Base window, right-click an article and select **Show Article**.

To move or copy an article:

1. In the Knowledge Base window, navigate to an article and drag and drop it into another group.
2. Select **Move** to move the article, **Copy** to make a separate copy of the article, or **Link** to create a copy of the article that is linked to the original article.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete an article:

1. In the Knowledge Base window, right-click an item and choose **Delete Article**.
2. In the dialog box, click **Delete**.

Managing Knowledge Base Article Groups

Knowledge Base article groups can be used to organize similar or related articles in a single location. For example, you could create a Denial of Service group to store specific articles about Denial of Service attacks such as a Ping Flood attack.

Groups and articles can be managed with drag and drop functionality. You can move or copy groups and articles into other groups within the Knowledge Base resource tree. If a group is deleted, the articles within that group are also deleted.



Note: To copy multiple resources at once, use Copy and Paste. You can drag and drop only one resource at a time.

Use the procedures in this topic to create, edit, move, and manage your KB groups.

Where: **Navigator > Resources > Knowledge Base**

To create a knowledge base group:

1. Right-click a group and select **New Group**.
A name text field appears under the group you selected.
2. In the name text field, enter a name.
3. Press **Enter**.

To rename a knowledge base group:

1. Right-click a group and select **Rename**.
2. In the name text field, rename the group.
3. Press **Enter**.

To edit a knowledge base group:

1. Right-click a group and select **Edit Group**.
2. In the Group Editor, type in the **Name** and **Description** text fields.
3. Click **OK**.

To move or copy a knowledge base group:

1. Navigate to a group and drag and drop it into another group.
2. Select **Move** to move the group, **Copy** to make a separate copy of the group, or **Link** to create a copy of the group that is linked to the original group.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete a knowledge base group:

1. Right-click a group and choose **Delete Group**.
2. Click **Yes**.

To updates to knowledge base groups and articles:

If a group or article name was modified through another ArcSight Console, or a group or article was deleted, renamed, or moved from another ArcSight Console, refresh shows those changes.

Right-click the **Knowledge Base** group or article and choose **Refresh**.

Associating Knowledge Base Articles

Knowledge Base groups and articles can be associated with other resources such as cases, reports, or filters.

Where: Navigator > Resources > Knowledge Base

To associate resources with KB groups or articles:

1. Locate a resource or resource group as the target, for example, a case or case group.
2. Right-click the resource and select **Knowledge Base> Associate with**.
3. In the Knowledge Base Article Selector, select an article to associate with the resource.
4. Click **OK**.

To associate grid elements with KB articles

1. In a Viewer panel grid view, right-click an event attribute and select **Knowledge Base>Associate>Cell/Row/Column with**.
2. In the Knowledge Base Article Selector, select an article to associate with the grid view's selected cell (data), row (event), or column (attribute).
3. Click **OK**.

Chapter 24: Finding Resources

Apart from visually navigating the resources in the Navigator panel, you can also find items in resource trees by searching or by locating them.

The search capability uses conventional query elements to search the entire set of system resources, returning a ranked list of qualifying items. Each user sees only those resources for which they have permission, regardless of the query. You can search for a string in All Resources or within a particular resource, for example, Cases.

Related topics:

- ["How Fields are Indexed" below](#)
- ["Using Text Search Syntax" on the next page](#)
- ["Using the Search Field on the Console Tool Bar" on page 708](#)
- ["Using the Search Result Columns" on page 710](#)
- ["Locating Resources on the Navigator Tree" on page 710](#)

How Fields are Indexed

ESM indexes fields at multiple levels, the lowest level covering a limited number of fields. The scope of the search increases as the index level goes higher.

Search index levels

Indexed fields	Index level	Affected resources
Index only these fields: name, id, uri, description, type, disabled, lockedBy	0	
Index fields in level 0 plus additional fields such as an asset's IP address or host name.	1	Asset
Index all attributes associated with a resource. This level should satisfy most user requirements for text searches.	2	All other resources except Asset



Caution: Levels 3 and 4 are also supported and include relationships between resources. Increasing the index level may allow you to search deeper. However, the size of the index will increase and performance will be slow.

Default search level settings:

The default search index levels are defined as follows:

```
search.index.level=2
search.index.level.Asset=1
```

You can change the level for all resources or specific resources, as described in the following procedure.

To customize the search index level:

1. Refer to the instructions on editing the `server.properties` file in the [ESM Administrator's Guide](#).
2. To change the default of 1 for the Asset resource, insert the statement for the setting in this format:
`search.index.level.Asset=number`
where *number* can be **0** or **2**. For example, if you use 2, this means Asset will be indexed in the same way as all resources.
3. To change the default for all resources other than Asset, add the setting in this format:
`search.index.level=number`
where *number* can be **0** or **1**.
4. To customize the level for specific resources, add one statement for each resource you want to customize. For example:
`search.index.level.Resource=number`
where *Resource* can be any resource, and *number* can be **0** or **1**.

Using Text Search Syntax

The search feature uses the Apache Lucene syntax. Lucene's query parser interprets the following as special characters:

+ - & & | ! () { } [] ^ " ~ * ? : \

Do not use * or ? at the start of your search string because this syntax returns nothing.

Do not use the wildcard for IPv6 address searches. Instead, enter the exact IPv6 address.

If your query string includes any of these special characters, escape them with a backward slash (\) for your query to work correctly. For example, if your search string includes (1+1):2, you write it as

```
\(1\+1\)\:2
```

However, if your query string *starts* with a special character other than * or ? which is not allowed, enclose the entire string in double quotes. For example, to search for this Resource ID:

```
^VVsOXg4BABCAIEuBhILMyg==
```

Enter

```
"^VVsOXg4BABCAIEuBhILMyg=="
```

in the query text field. You also use double quotes to enclose a phrase, such as

```
"keep them together"
```

For additional information about the Apache Lucene syntax, go to

https://lucene.apache.org/core/5_5_1/queryparser/org/apache/lucene/queryparser/classic/package-summary.html

and refer to the topic, Query Parser Syntax. Write your queries using the documented conventions.

Apache Lucene Syntax Relevant to Querying for Resources

Query Elements	Descriptions
Full or partial strings	Phrases, words, or partial words. Examples: <pre>"Attack Notification"</pre> <pre>notification</pre> <pre>notif</pre>
Wildcards	Question marks (?) for single-character substitutions and asterisks (*) for multi-character substitutions. - Do not use wildcards to start the search string. - Do not use wildcards to search for IPv6 addresses. Use the full IPv6 address instead. Examples: <pre>attack??</pre> <pre>name:"attack??"</pre> <pre>attack*</pre>
Boolean Operators	Use AND and OR to join strings. Examples: <pre>attack AND high AND compromise</pre> <pre>attack OR high OR compromise</pre>

Apache Lucene Syntax Relevant to Querying for Resources, continued

Query Elements	Descriptions
Fields	Resource field labels (grid view columns) followed by a colon, with the data expressed as plain strings, Boolean strings, quoted strings, or parenthetical expressions. Examples: <pre>type:datamonitor AND name:"event counts"</pre> <pre>name:"address space"</pre> <pre>name:(address+space)</pre> <pre>name:(address space)</pre>
Exclusion	Use NOT, the minus sign (-), and the exclamation point (!) to exclude strings. Examples: <pre>at???? -attack</pre> <pre>at???? NOTattack</pre> <pre>at???? !attack</pre> <pre>at???? !attack !type:File</pre>
Proximity	Extend data-field queries' scope with a proximity factor expressed as a numeral following a tilde (~). The numeral sets the maximum number of words allowed between the specified words in the resources found. Examples: <pre>name:("top events"~1)</pre> <pre>name:("top events"~2)</pre>
Fuzzy	Broaden query results with a relative letter-substitution factor expressed as a decimal fraction following a tilde (~). The values 0.0 to 0.9 apply, with the higher values increasing the substitutions made in the string. Examples: <pre>name:mssp~0.2</pre> <pre>name:mssp~0.0</pre>

Entering Values: Examples

Fields	Details
Dropdown fields	For dropdown fields that offer a list of values, enter the specific list item. For example, for a case's Ticket Type, enter Internal.
Dropdown fields with code and value pairs	For dropdown fields that offer a list of codes and their corresponding values, enter the code only. For example, for a case's Frequency, enter 0 to denote Never or Once.

Entering Values: Examples, continued

Fields	Details
Search narrowed to specific fields	To narrow your research to a resource's specific field, use the format <i>resource:fieldName=somevalue</i> For example: case:name=MyCase The fieldname must match the database column name. Column names follow the camelcase format. You can derive the fieldname from the field's label on the Console. For examples: • If the label is Name, the fieldname is name. • If the label is Ticket Type, the fieldname is ticketType. • If the label is Estimated Restore Time, the fieldname is estimatedRestoreTime.



Tip: Refer to ["How Fields are Indexed" on page 704](#) for information on how to fine tune your search index level.

Using the Search Field on the Console Tool Bar

To use the Search field:

1. In the Search field  on the Console toolbar, enter a name or phrase. Refer to ["Using Text Search Syntax" on page 705](#) for guidance on search syntax.
2. Click the **Find Resource** button ().

The search results are displayed in the Viewer.

- Single-click an item to display a preview of its definition in the Details pane on the Viewer.

Or

- Double-click an to open its definition in an Editor in the Inspect/Edit panel.

To limit a search to a particular resource type:


1. Click the **drop-down menu** tab on the Search field.
2. Select a resource type from the menus.

Notice that some resource types have sub-types from which you can choose. If you limit the search to a resource type, an icon of that type replaces the Search icon in the Search field.

For example, to search for a name or phrase only in Trends:

1. Select **Reports > Trends** from the Search drop-down.
2. Enter the search string.
3. Click the **Find Resource** button.

The Search field in the toolbar accepts all the Query Options described in the following figure.



As an alternative to using the quick Search field option, you can get a full Search panel in the Viewer:

1. Select **Edit > Find Resource** in the Console's menus, or press **Ctrl+F**.
2. In the Viewer panel's Resource Search tab, enter a query string in the **Search query** line, set the number of results to allow, and click **Find**. See ["Using Text Search Syntax" on page 705](#) for guidance on search syntax.
3. On the returned results, click any item to see its details or click a result column heading to change the order.

When you click a resource listing in the **Details** panel, it shows you the various pieces of related system information that justified that item's ranking.

Using the Search Result Columns

The Find Resources viewer displays the resources found by the search. Click any column heading to toggle between descending and ascending order.

Column	Description
Score	Ranking of resources a query returns, based how frequently the search term appears in each resource.
Type	Top-level categorization of the resource as shown on the Navigator panel, for example, Active Channel, Asset, Rule, and so on.
Name	The full name of the individual resource.
URI	Full uniform resource identifier for the individual resource.

Locating Resources on the Navigator Tree

These steps provide a way to browse from the resource editor to the resource's exact location on the Navigator panel.

1. In an entry in a resources grid view, or in the top tab of a resource editor, right-click and select **Find <resource type> in Navigator**.
2. Look for the highlighted item in the Navigator panel's resource tree.

Chapter 25: Managing Resources

This chapter discusses the administrator tasks to manage ArcSight ESM.

Working with Resource Groups

You can group resource types in the Navigator panel to help you organize and manage them. Groups can also be hierarchical, resulting in resource trees in the Navigator panel. Apart from the characteristics of the resources involved, such as assets or vulnerabilities, each group identity has certain properties you can edit in the Group Editor.

Adding or Editing a Resource Group

To add or edit a resource group:

1. To add a group, right-click a resource group and choose **New Group**.
Or to edit an existing group, right click the group and choose **Edit Group**.

2. In the Group Editor, enter or change the group attributes you want to change.

For the group's name, do not use special characters like & (ampersand), * (asterisk), and % (percent). The group name becomes part of URIs (uniform resource identifiers), and therefore follows the same restrictions for URIs in general. The escape character option is not available in this case.

Optional: Under the **Reference Pages** section, enter applicable information, for example:

- In **Group Page**, enter the URL to the website (for example, a wiki page) containing information relevant to the resource group.
- In **Members Page**, enter the URL to the website (for example, an internal site) listing team members associated with this resource group.

Entering data in the Common and Assign sections is optional, depending on how your environment is configured. For information about the Common and Assign attributes sections, as well as the read-only attribute fields in Parent Groups and Creation Information, see ["Common Resource Attribute Fields" on page 725](#).

3. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
4. Click **Apply** to put your changes into effect but leave the editor open. Click **OK** to apply your changes and also close the editor.

Fields containing system information (like Creation Time) are not editable.

See ["Reference Pages" on page 1027](#) for more about using the **Group Page** and **Group Children's Page** fields.

See ["Job Scheduler" on page 1003](#) for information about scheduling tasks or “jobs” for reports (individually or by group), rules, or Threat Detector snapshots.

Using the Categories Tab for Asset Groups

Where: Navigator > Resources > Assets > *asset group*

Right-click your asset group and select **Edit Group** to display the group editor..

The asset group editor for has an additional Categories tab. This tab has two sub panels: **Local Asset Categories** and **Inherited Asset Categories**. **Local** shows assets that are explicitly assigned to categories. **Inherited** shows assets whose category connections are presumptions based on a parent's group or a simple asset-range association.

Moving, Copying, Linking, and Deleting Resources

You may need to move or duplicate a resource to better organize your work or to make editable copies. You may also need to delete resource definitions you no longer need. These tasks are described here.

Where: Navigator > Resources <resource>

To move, copy, or link a resource:

1. Select the resource type you want to work with (Active Channels, Filters, Rules, and so on).
2. Drag and drop it into another group of the same resource type. The system displays a popup that provides options to move, copy, or link the resource.

Select **Move** to move the resource, **Copy** to make a separate copy of it, or **Link** to create a copy of the resource that is linked to the original.

If you select **Move**, the resource or resource group moves to the new location. If you select **Copy**, you create a separate copy of the resource or resource group that will not be affected when the original resource or resource group is edited. If you select **Link**, you create a copy that is linked to the original resource or resource group. Therefore, if you edit a linked item, whether the original or the copy, all links are also edited. When deleting linked items, you can either delete the selected item or all linked items.

To delete a resource:

1. Select a resource.
2. Right-click and select **Delete <resource>**.

Locking and Unlocking Resources

The locking and unlocking capability applies to the following ArcSight content:

- System core content
- User created content

ArcSight Standard Content

A set of predefined standard content is installed by default. This content provides the foundation building blocks for ESM to work.

Standard content is available in the ArcSight System sub-tree of each resource tree. For example, core content for the Filters resource is available in Shared/All Filters/ArcSight System/.

The modification of ArcSight System content can adversely impact operation, therefore, it is locked by default. Open Text strongly recommends against unlocking or modifying this content. To unlock this content, contact Customer Support.



Note: Use the resources available in ArcSight Administration and other content downloaded from the Marketplace to create content that suits your needs. Make copies of these resources in your own personal folders and customize there.

User-created content

ArcSight users can lock any resource or a group of resources to which they have write access privileges. Locking prevents a resource from being deleted. Once locked, such resources or groups can be unlocked only by these users:

- The lock owner, meaning, the user who locked the resource.
- Any user who has write permissions to the lock owner. That is, a user who has privileges over the user who applied the lock. For example, the administrator user has write permissions over all users by default. Therefore, if user joe locks a resource, the user administrator can unlock it.



Tip: You can make a copy of a locked resource even if you do not have the privileges to unlock it.

You can edit resources in a locked group if you have write access privileges to the resource, however you cannot do the following:

- Delete or remove resources from it.
- Add a new resource to it.

To unlock a resource:

Right-click the locked resource and select **Unlock** from the drop-down menu.

Selecting Resources

You often need to select resources to act on or use while authoring or configuring analysis tools. Selecting is often the first step in managing, authoring, or analyzing resources.

While the Navigator panel is your usual means of selecting resources, you can also encounter the Select Resources dialog box any time selection is a necessary part of some task, such as adding a case group to a rule action or adding user groups to access control lists (ACLs).

For resource groups, click to select the group you want to choose, then click **OK**. For options that allow multiple selections, select the check boxes next to individual entries in the list under a group, then click **OK**.

This dialog is also displayed for setting user permissions on resources and operations.

More information:

- For information about setting permissions on resources, see ["Managing Permissions" on page 107](#).
- For information about setting action permissions on who can deploy data monitors, see also ["Controlling Who Has Permissions to Deploy Data Monitors" on page 116](#).

Visualizing Resources

The resources presented in the Navigator panel or graphically in the Viewer panel are organized into hierarchical groups for easy browsing. Among similar types of resources, there can be logical relationships. Graphs can make these relationships readily visible.


Graphing Resources

Where: Navigator > Resources

Procedure:

1. Right-click one or more individual resources or resource groups except Notifications.
2. Select **Graph View** in the context menu.

The Viewer panel graphs the resources in a new channel. The following example shows a custom rule depicted in the center.



From the graph, you can view its relationships specified in the rule's Conditions tab. As the example shows, the rule interacts with two lists and two filters.

Using Graphs

Once generated, you can manipulate graphs further. There is a set of command buttons at the top of the view and a parallel set of commands available by right-clicking the graph itself.



Note: The Zoom In, Zoom Out, and Fit Content options do not work with event graphs created after version 7.0. Scrolling works but can cause issues with large filter sets. To use these options with event graphs, you must enable the **Use classic charts** option in **Global Preferences**.

Resource Graph Command Buttons

Command	Button	Description
Fit Content		Sizes the graphic to the available display space.
Zoom In / Zoom Out		Increases or decreases the size of the displayed graphic.
Zoom Selected		Zooms in on a selected portion of a graphic.
Hierarchic Layout		Presents nodes in a vertically descending cascade, similar to a family tree. Hierarchic layouts are appropriate when viewing event relationships that have a common root.
Organic Layout		Displays nodes in an arrangement based on minimum edge length, which tends to cluster nodes that relate to a common node. Likewise, node clusters with nodes in common also tend to group together.
Circular Layout		Positions nodes in hub-and-spoke arrangements with each node radiating edges to, or receiving edges from, the nodes with which it interacts. Circular layouts are most useful when multiple roots are present or there are a number of source-target relationships to clarify. If an organic layout is difficult to read because the edges are too dense, try a circular layout instead.
Orthogonal Layout		Arranges nodes on the basis of logical connections, using electrical schematic-style right-angle layouts. These layouts are very useful for clearly tracing connections and identifying node clusters.
Overview		Opens a reduced rendering of the entire graph. You can drag the highlighted section in the reduction to move the displayed area in the main view.
Hierarchy Tree		Opens a complete list of the nodes plotted in graphic layouts. Click a node in the list to scroll to that node in the main view.

Resource Graph Command Buttons, continued

Command	Button	Description
Print		Prints the displayed graphic.
Add Graph View to Case		Adds the current graph view to a case you select. Choosing this option opens the Case Selector dialog, where you can browse cases. Select a case to which to add the current graph view and click OK on the Case Selector dialog. The graph view is added to the selected case as an attachment, accessible on the Attachments tab in the case editor for that case.
Help		Display the relevant ArcSight Console online Help topic.

Configuring Resource Graphs

Where: Navigator > Resources

Procedure:

1. Right-click one or more individual resources or resource groups except Notifications.
2. Select **Graph View**.
3. Click anywhere in the resulting graph, right-click, and select **Configure Resource Graph**.
This opens the Configure Resource Graph dialog where you can specify which resources to display in graph views.
4. Select resources to show or hide.
5. Click **OK**.

Viewing Resources in Grids

While the grids you see in the Viewer panel are most often views of events, these grids can also display organized sets of information about resources in the Navigator panel.

In the Navigator panel, certain resource groups include **Grid View** in their right-click context menus. This command causes the items in the group to display in a grid view, where you can:

- Review them using the sorting and column customization features that grid views offer.
- Right-click resource items in grid views and use the same context commands that those resources have in the Navigator panel.

Validating Resources

Resources can break or become invalid because they are improperly built or cannot find other resources they depend on.

Resource validation takes place automatically during an upgrade, package import or export, or when you insert or update a resource. This topic explains how to manually identify, troubleshoot, and fix broken resources.

About Valid and Invalid Resources

Valid resources appear in the Navigator with their associated icons. For example, the Navigator displays a valid filter like this: .

A valid resource is fully available to other resources that reference it. The resource can participate in the event flow, trends, reports, data monitors, channels, filters, rules, and so forth.

A resource can "break" or become "invalid" either because it is constructed improperly. For example, an active list schema does not match the underlying table. Changes to resources can also break other resources that have dependencies. For example, a filter referenced by a rule is no longer available. This happens when a resource being depended on is deleted, renamed, or moved in the Manager; or not retained during an upgrade, import, or export.

Invalid resources appear in the Navigator as broken or torn. For example, and the Navigator displays an invalid filter like this: .

An invalid resource includes an [Invalid Reason](#) field under on the Attributes tab of its editor, as described in.

An invalid resource cannot participate in the event flow or other resources in real time. For example, an invalid asset cannot participate in event asset resolution. Correlated events in which the source or target address points to the invalid asset are not generated. Similarly, an invalid rule does not trigger and generate correlation events.

Fixing and Validating Resources

When a resource become becomes invalid, its Editor includes a **Validate** button that you can use to test and validate the resource after you fix it. Clicking the **Validate** button on a resource that was previously broken results in a check of the resource logic and dependencies. If the system determines the resource is now valid, the resource icon in the Navigator is updated to

reflect a working resource. If the system determines the resource is still broken, an error message describes the problem.

High level steps to fix and validate a resource:

1. Identify an invalid resource. Sometimes problems with dependencies, for example filters or rules which are used in many other resources. are a result of broken resources.

A valid resource looks like this: , and an invalid resource looks like this: 

For example, if "My Top Threats" filter depends on "My Hotlist" filter, removing or renaming "My Hotlist" filter breaks "My Top Threats" filter.

A scheduled job (like a scheduled rule group or archived report) can also break if one of the resources it depends on is missing. The broken icon for a scheduled job shows on the Current Jobs list.

2. If you do not already know why a resource is broken, double-click the resource in the Navigator panel and click **Validate** in the resource editor.

Validation gives you an error message that describes the problem, for example, this resource depends upon another resource that is now missing. The error dialog includes a Copy option for copying longer messages to an external editor.

3. Fix the problems with the resource. See ["Troubleshooting Requirements for Valid Resources" on page 721](#).

To continue with our example, adding back in the filter "My Hotlist" would fix the problem we mentioned in the beginning of the procedure.

4. In the resource editor, click **Apply** to save changes to the resources you modified.



Tip: For problems that can be validated on the local client, you can click **Validate** before clicking **Apply**. If the resource is fixed, its "working" icon is immediately reflected in the Navigator. However, for other types of problems; you need to **Apply** the changes to the resource before you **Validate** the resource. This is because some types of changes must be processed on the Manager. Dependencies and relationships to other data may not be available on the Console client.


If you think you have fixed a resource but it is still not showing as fixed in the Navigator, make sure you **Apply** all the changes you made to it and then click **Validate** again.

5. In the resource editor for the resource that was broken, click **Validate** again. If the

resource passes validation, its icon in the Navigator updates to reflect a working resource.

In the resource Editor for the resource that was broken, click the **Validate** button. If the resource passes validation, its icon in the Navigator updates to reflect a working resource. Otherwise, the broken icon remains and an error message describes the problems.

Some problems require saving fixes to the Manager, so be sure to click **Apply** and save changes to resources you fix before you click **Validate**.



To validate a scheduled job, click the **Open scheduled jobs list** tool button () to display scheduled jobs in the Viewer, right-click the job you want to validate, and choose **Validate** from the context menu. If the job passes validation, its icon in the Current Jobs list updates to reflect a valid task.



Note: ESM administrators can use a command line utility, `resvalidate`, to validate resources. Refer to the [ESM Administrator's Guide](#) for a description of the command. Running this command is part of the ESM upgrade process. Refer to the [ESM Upgrade Guide](#) for information.

Troubleshooting Requirements for Valid Resources



Caution:

It is possible that dependent resources are pointing to the wrong resource. This usually happens if you rename a resource, then re-use the old name on a new resource of the same type. The dependent resources will be linked to the old name. To avoid this problem, don't re-use an old name on resources of the same type. An example of a dependent relationship is that of a query depending on a trend.

The most common cause of an invalid resource is a dependency issue; another resource that the broken resource depends on is missing from the database. Some resources have additional requirements or limits that can also affect validity. Following is a summary of requirements for creating valid resources.

If any of these requirements are not met, the resource will break. To fix the resource, edit its definition to be in line with these requirements.

- **All Resources** - If the definition for a resource references another resource, the referenced resource must be available in the Manager database. This requirement is true for all types of resources.
- **Devices and Assets** - Each asset address must be unique within a zone, an asset can belong to one zone only, and the asset IP address must fall within the address range of its network zone.
- **Device and Asset Ranges** - Start addresses must be less than end addresses, asset ranges must be within the address range of the associated network zone, and asset ranges should not overlap another asset range in the same zone.
- **Zones** - Start addresses must be less than end addresses and network zones should not overlap other zones in the same network.
- **Reports** - Report templates cannot contain more than 20 charts or more than 15 tables.
- **Active Lists** - Active List schema must match the underlying table and must not include programming errors.

The following table lists the resources that can become invalid:

Reasons for Invalidated Resources

This resource becomes invalid...	when it violates one or more of the following constraints...	which results in...
Device/Asset	- Asset address must be unique within a zone. - An asset only belongs to one zone. - Asset IP address must fall in the address range of its network zone.	The invalid device/asset cannot participate in the event asset resolution. Therefore, if an event source/target address points to the invalid device it cannot be resolved.
Device/Asset Range	- Start address must be less than end address. - Asset range must be within the address range of its network zone. - Asset range should not overlap another asset range in the same zone.	The invalid device/asset range cannot participate in the event asset resolution. Therefore, if an event has its source/target address fall in an invalid device range its asset resolution cannot be resolved.
Zone	- Start address must be less than end address. - Network zone should not overlap other zones in the same network.	The assets falling within this invalid zone get invalidated and cannot participate in the event asset resolution.
Filter	Dependency constraint. For example, a filter may depend on other resources, like asset, active list, vulnerability etc.	The invalid filter causes the resources that depend on it to get invalidated.
Rule	Dependency constraint. For example, a rule may depend on other resources, like filter, asset, vulnerability, active list, session list etc.	The invalid rule cannot be triggered, so the corresponding correlation events are missed.
Data Monitor	Dependency constraint. For example, a data monitor may depend on other resources such as a filter.	The invalid data monitor stops fetching live data to feed the dashboard.
Active Channel	Dependency constraint. For example, an active channel may depend on other resources such as a filter, or asset vulnerability.	You cannot attach or open an invalid active channel
Report	Dependency constraint. For example, a report may depend on other resources, such as filter or asset, vulnerability and active list.	You cannot run the invalid report manually from console or as a scheduled task.
Trend	Dependency constraint. For example, a trend that depends on a query is invalid as soon as a query is changed.	The invalid trend stops generating any trend data.
Scheduled Task	Dependency constraint. For example, a scheduled task may depend on other resources, such as filter.	The invalid scheduled task cannot run.

Reasons for Invalidated Resources, continued

This resource becomes invalid...	when it violates one or more of the following constraints...	which results in...
Report Template	The report template cannot contain more than 20 charts or more than 15 tables.	The invalid template causes the reports that depend on it to be invalid.
Profile	Dependency constraint. The Profile depends on resources such as the filter it uses to determine which events to run discovery on. It also depends on the group where snapshots and patterns are saved. All these resource must exist and the creator should have appropriate permissions for them.	This resource is invalidated and the scheduled runs may be skipped.
Active List	If the Active List schema does not match the underlying table etc, or due to some programming error.	The resources (Rules, reports etc.) that are dependent on the Active List get invalidated
Focused Report	Dependency constraint. For example, a focused report may depend on other resources, such as a report, filter or asset.	The invalid focused report cannot be run either manually from the Console or as a scheduled task.
Query	Dependency constraint. For example, a query may depend on other resources, such as a filter, asset, or active list.	The invalid query causes the resources that depend on it, such as report and trend, to become invalid.

Resource Validation During Upgrade or Package Import

If the Manager detects a conflict during an upgrade or import process, it invalidates the conflicting resource, and continues with the upgrade or import process. The dependent resources for the conflicting resource is automatically re-validated and disabled after the resource validation process completes.

- After an upgrade process, a report called `validationReport.html` is generated in the `<ARCSIGHT_HOME>/upgrade/out/<time-stamp>` directory. See the [ESM Upgrade Guide](#) for more details.
- After an import process, you can check the Console to make sure that you do not have any invalid resources. You are expected to fix the invalid resources manually.

After you resolve the conflict, the dependent resources for the conflicting resource is automatically re-validated.

Extending Audit Event Logging

Updates to existing resources are logged as audit events, as described in ["Audit Events Common to Most Resources" on page 834](#).

If you want to get additional details within the audit events on resource updates beyond what is provided by default, you can enable a resource audit property on the Manager to specify which resources should show extended audit event information.

To configure resources for more detailed update auditing, add a URI to the `resource.audit.update.uris` property in the `server.properties` file. For example:

```
resource.audit.update.uris=/All Users/
```

turns on extended audit logging for all resources under the `/All Users/` subtree.

Leaving this property blank would turn this feature off (and show only default audit information).

To show detailed audit information for multiple resource types, list resource URIs separated by commas (no spaces). For example, to show extended update audit logging for users and system assets, set the property like this:

```
resource.audit.update.uris=/All Users/,/All Assets/ArcSight System Administration/
```

Refer to the topic, "Managing and Changing Properties File Settings," in the [ESM Administrator's Guide](#) for instructions.

Extended information on the resource update is logged in two places:

- In the internal audit event generated for the resource update, `Device Custom String5` is set with the update information. The audit event information is shown in the `Device Custom String5` field in this format:

```
<UUID generated for this change>:[<name of attribute>:<old value>:<new value>]+
```

- The update information is also written to a log file, `<ARCSIGHT_HOME>/logs/default/resource_update_audit.log` file. The audit event information is shown in the log in this format:

```
<UUID generated for this change>:<URI of resource>:<ID of resource>:[<name of attribute>:<old value>:<new value>]+
```




Tip: How to interpret the resource update log:

- The “+” in the message format examples above is regular expression notation used to indicate that there can be one or more of *<name of attribute>:<old value>:<new value>* triplets shown in the audit event.
- Any “:” character in any attribute name or value is escaped with a backslash to “\:”.
- Any “\” character in any attribute name or value is escaped with a backslash to “\\”.

Common Resource Attribute Fields

The following fields are common to several types of resources. You can find these fields in the resource editor Attributes tabs under the Common, Assign, Parent Groups, Creation Information, and Last Update Information sections. (See also ["Resource Attributes" on page 1032.](#))

Common

Entering data in the **Common** section is optional, depending your environment setup.

Common Fields	Description
Resource ID	Read-only field that shows the ID that ESM has assigned to this resource when it was created.
External ID	An identification string suitable for, and which can be referenced by, systems outside ESM. Common applications of External IDs include appropriate naming for Case and Asset resources that are tracked in common with defect reporting or vulnerability-management systems. If your system interfaces with a third-party incident tracking system, such as Remedy, enter an ID that corresponds to that system. Your administrator can advise you on the correct values for this field, if applicable.
Alias (Display Name)	An optional alternate identification string used for referencing resources. If given, this alias appears in place of the resource's name everywhere it may be seen. Your administrator can advise you on the correct values for this field, if applicable. If you use an alternate event naming scheme in your environment, enter an alias for this resource here.
Invalid Reason	If a resource is broken or invalid, an “Invalid Reason” field is included in its Attributes table. An abbreviated explanation is shown in this field. (See also "Validating Resources" on page 718.) Click the browse button  at the end of this field to get a popup dialog that shows the full text of the explanation.

Common Fields	Description
Description	Description of the resource. Use this field to communicate the purpose of this resource to other users. For example, if this is a resource that leverages or depends on another resource (for example, a query viewer or trend that uses an SQL query), this is a good place to make note of that relationship.
Version ID	The globally unique version ID for this resource. Version IDs are assigned when you export a resource as part of a package, if the resource has changed.
Deprecated	Toggle to indicate whether the resource is current or deprecated (obsolete).

Assign

Assign Fields	Description
Owner	A user selected from the Users resource tree who should be notified about this resource.
Notification Groups	The user groups selected from the Users resource tree who should be notified about this resource.

Saving Copies of Read-Only Resources

Although you may be limited to read-only access to certain resources in the Navigator panel, you do have the option to save a copy of such a resource to your own group where you have write access.

Procedure:

1. Click the **Save As** button to make a copy of the resource and save it in a specified group to display the resource group selector dialog.
2. Select the group in which you want to save a copy of the resource.
3. Specify the name you want to assign to your copy of the resource.
4. Click **OK**.

The resource copy appears in the resource tree. You have write permission on this copy.

The Connectors, Users, and Notification editors do not support **Save As** functionality. In these editors, you see the **OK/Cancel/Apply** buttons, but the fields for those resources are read-only.

Managing File Resources

The Files resource tree, when populated, lists various files that have been saved as resources so that they are accessible to all users of the system who are authorized for such access. File resources include case file attachments, templates, and general-purpose shared files.



Note: Case attachments are stored twice, at a minimum: in the case itself, and in the File resource. The case owner can additionally share the case attachment with a custom file resource group (in addition to the default /Attachments group). Case attachments therefore use up at least twice the storage size. See ["Best Practices on Attaching Files to a Case" on page 649](#) for information on managing file sizes.

In addition to the tasks in this topic, you can also rename or lock a file, get a Graph View of a file, and so forth. Simply select the file in the Navigator, right-click, and choose a menu option. Operations on groups are also available. Options may vary depending on which file or folder you have selected in the Navigator.

Related topics:

- ["Uploading Files and Creating a File Resource" below](#)
- ["Working with Files" on the next page](#)

Uploading Files and Creating a File Resource

Where: Navigator > Resources > Files

Procedure

1. Right-click a file group and select **New File**.
2. At the bottom of the File Editor panel, click the **Upload** button and select the local file to add.
3. On the Upload File Content browser, select the file to and click **OK**.
4. On the File Editor Attributes tab, enter values for the attributes that identify the file.

The Name attribute is initially the same as the Filename attribute, but you can change the Name.

Certain attributes are read-only: Upload type is Console. Filename, File size, and Mime type are set based on the selected file.

5. Click **OK**.

Working with Files

Where: Navigator > Resources > Files

To view files:

- Right-click a file resource and select **Open**.

Or

- In the File Editor, click the **Open** button.

The file is downloaded to a temporary subdirectory called `arcsight-files` and is displayed in a web browser.



Note: File resources can be downloaded as often as needed by any Console user authorized to access the file resources. Downloading a file does not change the file resource, or the shared file contents on the server.



Note: Case attachments and Manager-generated files are stored in `Files/Shared/All Files/Attachments`.

To download files locally:

1. Right-click a file and select **Download**.
2. Specify a location and file name for the new local file.

You can also Download a file resource from the File Editor by clicking the Download button.

To edit file resource attributes:

1. Right-click a file and select **Edit File**.
2. Change the values, as appropriate.
3. Click **OK**.

To replace file resource contents:

1. Right-click a file and choose **Edit File**.
2. In the editor, click **Replace** and select the local file containing the new contents for the file resource. The file resource name changes if the selected local file has a different name.
3. Click **OK**.

To delete file resources:

1. Right-click a file and select **Delete File**.
2. Click **Yes** to confirm.

To add a folder or file to a Package:

From the Files resource Navigator, you can add a file or folder to an existing package or create a new package and add the file to it.

1. Right-click a file or folder and select **Add to Package**.

This opens the Package Selector.

2. In the Package Selector, do one of the following:

- Navigate to a package to which you want to add the file or folder, and click **OK**. The file is saved to the selected package.
Or
- Navigate to a location where you want to create a new package and click **New Package**. This opens the Package Editor where you can name and configure the new package. The selected file or folder is included in the new package.

See ["Managing Packages" on page 730](#) for more information..


To find a file:

- To find files stored on the Manager, select **Files** in the Navigator and browse the folders.
Or
- Select **Edit > Search** from the menu, enter a file name in the **Search query** field, and click **Find**.

See ["Finding Resources" on page 704](#) for more information.

Chapter 26: Managing Packages

Packages are collections of resources that can be installed into the system resource tree.



To access available packages, click the **Packages** tab of the Navigator panel. The package tree appears, and you can click to expand installed packages to see the resources within. When the package button in the upper left is highlighted with a dark background () it displays uninstalled packages: this is the **Advanced** view. In the Advanced view, it also shows all the resources on which explicitly included resources depend. For example, a report might be explicitly included in a package, but any queries and templates are also implicitly included, unless they are explicitly, directly or indirectly (think groups), excluded.

The icon for uninstalled packages is grayed out (). When you click this button, the Normal view does not show uninstalled packages. The Normal view of the package contents only shows resources explicitly included in the package. When you toggle between views, each view remembers which packages were expanded, and restores the tree to that state.

Right-click options for viewing package information

These options list all resources in the package, including details such as type, full path to location in the tree, and resource name.

Option	Description
Show Package Archive Contents	In the Advanced view, you can right-click an uninstalled package and select this option to show the resources in the package. The result indicates resources that have been deprecated, if any. The Resource Count in the header includes resources that are not exposed to the user.
Show Current Package Contents	In either view you can right-click an installed package and select this option to show the resources in the package. The result indicates resources that have been deprecated, if any. The Resource Count in the header includes resources that are not exposed to the user.
Compare Archive with Current Package Contents	Right-click an installed package and select this option to show the resources in the package. The comparison indicates what type of change was done to a specific resource type. The result includes only the resources that are visible to the user, so the Resource Count may be less than the count shown in the other viewing options.

Creating or Editing Packages

When you add a group to a package, all the group's contents are automatically included. For top level groups (<Resources>/All <Resources>), the package therefore includes all folders that come with ArcSight, which may present a problem if the package is imported in another ArcSight system. If you defined resources directly under the All <Resources> node and you want to add these resources to a package, create a group and add the resources there.

It is also important to note that when you delete a package containing a group, you delete members of the group that were added after the package was created. If it is a top-level group (<Resources>/All <Resources>), you would be deleting all of those resources. Packaging sub groups gives you more flexibility and less risk. Also see ["Backing Up and Restoring with Packages" on page 740](#).

When editing packages, you should consider whether you need to export it before you change it so you have a backup, and possibly export it when you are done, to update or create a new backup.



Tip: Organizing your package contents

Users typically create packages according to resources such as a package for lists, another package for queries, and so on.

You can also consider creating packages according to related resources comprising a specific use case. For example, if you have created rules that act on specific lists, include those resources in the same package.

Where: Navigator > Packages tab

To create or edit a package:

1. To create a package, right-click a package group and select **New Package**.
To edit a package, right-click the package and select **Edit Package**.
2. In the Inspect/Edit panel, set or change the attributes as appropriate. Some of these attributes may not appear if they do not apply. The only required field is the **Name** field and many field values are supplied by the system:

Package Editor Attributes
Name Enter a name for the new package.
Required Packages Specify the packages that must be installed for this package to function.
Optional Packages Specify packages that are related to this package, but which are not required for it to function.
Required Features Enter any features that must be available for this package to function. The list contains Threat Detector.
Installed This is a read-only status indicator. If it is checked, it indicates that the package is installed, which new packages are, by default. You can uninstall a package by right-clicking it in the tree after you have created it, and selecting Uninstall Package. The icon for uninstalled packages is grayed out.
Update Available This is a read-only status indicator. If checked, it indicates that a different version of the package was imported, but has not been installed. The icon for packages in this state has a small white up arrow in the lower left corner ().
Author Name Enter the name of the author or source for the new package.
Package Version The package version can be any string, but the recommended format is <i>n.n.n.n</i> , with numbers in decreasing importance (major, minor, release, build).
ArcSight Version This is a read-only ESM version that is the minimum product version needed to support this package.

Package Editor Attributes

Format

Specifies the format to use for archiving. This affects the resources that are included in the package, the resource attributes that are retained in the package, and how conflicts are handled during package installation.

Tip: If your package contains lists, see also ["List Data" on page 741](#) for additional help with formats you can use.

- **Contentsync** - Use this if you intend to synchronize content among other Manager peers. The Manager that is the source of the contentsync package is the designated publisher, while other Managers are the subscribers (content management feature). For details on content management and the establishment of peer relationships, refer to the [ArcSight Command Center Users Guide](#).

Note: Not all packages are eligible for content synchronization. See ["Supported Packages for Content Synchronization" on page 735](#) for more information.

- **Default** - Use this for backing up resources on a Manager. This format includes additional, such as data in active lists, including information specific to a Manager, whereas other options do not.
- **Export** - Use this for packaging resources for transport between systems. Manager-specific information is excluded from the exported package for resources with attributes that would otherwise retain such information upon a "Default" export.

For example, a trend packaged in "Export" format does not include Start End Time trend attributes nor original table IDs. Instead, the imported trend uses Start and End times corresponding to when the package is installed on a new system. Also at install, a new trend table is created. (See descriptions of Imported Trend Start Time and Imported Trend End Time fields under ["Trend Attributes" on page 471](#).)

Similarly, active lists and session lists packaged in "export" format do not include Locked By attributes, table IDs, or session/active list entries. New tables are created when the lists are imported, and the other attributes are tracked when these resources launch on the new system.

This format packages other resources similarly as a means of optimizing portability for content distribution.

Standard system content is packaged using the "Export" format. Also, Managed Security Service Providers (MSSPs) who provide content to installations at various customer sites might package resources in this format.

- **Exportuser** - Use this format only for exporting user accounts with no permissions, personal group information, or relationships to other resources. If you want to export user accounts that include permissions and groups, use the default format instead.
- **Upgrade** - For use by ArcSight Professional Services only. This format might be used for resource upgrades of older systems in some circumstances. (Usually, standard upgrade utilities and processes are used instead.)

Obfuscated

Check this box to encrypt the contents of the ARB file, making it impossible to view without importing it.

Package Editor Attributes

Exclude Reference IDs

Check this box to remove reference IDs from the package when it is exported. Generally, you would exclude reference IDs only when you plan to import the package into a different system. Leave the box unchecked to include reference IDs, which improves performance for packages that are imported to the same Manager from which they were exported.

Note: Reference IDs are not the same as resource IDs. Resource IDs are always part of the resources and included in package imports and exports.

Creation Timestamp

The date and time when this package was last exported (archived).

Required package For

Shows any other packages for which this package has been set as a required package. These other packages are therefore "Dependent Packages."

Optional Package For

Shows any other packages for which this package has been set as an optional package.

Archive ID

This system-assigned ID that is refreshed whenever you export the resource.

Available Archive Translations

If the archive defining the package has been rendered into different languages, they are listed here.

Also refer to the common attribute fields described in "[Common Resource Attribute Fields](#)" on page 725.

3. Click the **Resources** tab in the Package Editor.
4. Click the **Add** drop-down menu to select the resources that this package should contain. You can select groups or individual resources. Check the Children Only box to include resources below the specified resource in the tree. For example, selecting the group /All Session Lists/ArcSight Administration/User and checking **Children Only** would include only the session list resources in that group, not the group itself.
5. Check the **Only If Referenced** box to conditionally include resources if they are referenced by other resources. This is best when used in conjunction with the **Children Only** attribute, and you are adding a group resource.
6. To exclude resources from the new package, use the Removed Resources panel in the lower half of the Resources tab.
 - To exclude resources by type, click the **Excluded Resource Types** tab and select from the list of available types.
 - To exclude specific resources, click the **Removed Resources** tab, click the **Add** drop-down menu, then select the resources you want to exclude. This tab also includes a

Children Only option, but an **If Not Included** option instead of an **Only if Referenced** option.



Caution: The only way to exclude Asset Category resources from a package is to specify the Asset Categories explicitly using the Removed Resources tab.



Tip: If you include locked resources or shared system resources, then when uninstalling or deleting a package, you get a message that the package framework has a number of locked resources, and therefore cannot be uninstalled or deleted.

When that happens you can either unlock the resource and continue or choose to skip the locked resource, in which case it will not be uninstalled or deleted.

7. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).

About Locked Packages

You can lock packages after creating them by right-clicking on the Navigator panel and selecting **Lock Package**. Note that if you send locked packages, these packages cannot be installed in another ESM system.

Adding Resources from the Resource Navigator

You can add to a resource to an existing package by using the right-click menu on a selected resource in the Navigator tree.



Tip: Do not add shared system resources and locked resources to packages. If you do, uninstalling such packages will fail.

Where: **Navigator > Resources > resource**

1. Right-click the particular resource you want to add (for example, My Report).
2. Select **Add to Package**. The system displays dialog.
3. In the Package Selector dialog, select a package to which to add the selected resource and click **OK**.

Supported Packages for Content Synchronization

Content synchronization provides the ability to create content in one ESM installation and push this content to other ESM installations. With the use of the content management feature, you

can establish peer relationships and designate the publisher and subscribers. You can then manage packages to be pushed automatically or manually. Details on content management are in the [ArcSight Command Center Users Guide](#).

Not all packaged resources are eligible for content synchronization. The following table contains a list of resource packages you can push to subscribers.



Note: It is assumed you have created the package with the contentsync format, as described in ["Creating or Editing Packages" on page 731](#).

Resources Eligible for Content Synchronization

Resource	Notes
Active channels	
Active lists	Definition only, not list entries. Subscribers can override with their own entries.
Drilldowns	
Drilldown lists	
Dashboards	
Data Monitors	Definition only, but not data. Subscribers can override with their own data.
Customers	
Fields	
Field Sets	
Integration Configurations	
Integration Commands	
Integration Targets	
Knowledge Bases	
Notes	Notes are defined within resources.
Queries	
Query Viewers	
Reports	
Report Templates	
Rules	
Saved Searches	Created on the Command Center and available on the ArcSight Console as resources.
Search Filters	Created on the Command Center and available on the ArcSight Console as resources.

Resources Eligible for Content Synchronization, continued

Resource	Notes
Session Lists	Definition only, but not list entries. Subscribers can override with their own entries.
Trends	Definition only, but not the events searched by the trends.
Users	

Exporting Packages

Exporting enables you to save a package as a file with the .arb extension in the folder of your choice. You can use .arb files for backup or to transport them to other systems.

Where: Navigator > Packages tab

Procedure:

1. Select one or more packages to export.
2. Right-click and select **Export Package to Bundle**.

You can also export after editing a package by clicking the **Export** button at the bottom of the edit panel.

3. Enter a file name and folder for the file. Leave the default extension as .arb.

The console generates a message that indicates whether objects are missing from the export. If objects are missing, the console lists them. If no objects are missing, the list is blank. If the list is *not* blank, it is possible that some of your selections could result in information being lost upon import. For example, if there are multiple rules in a group and you select to export the rules but not the group, when you import the package to another system, the group is created but no other information is included. The message is intended to verify that you meant to exclude certain objects and allow you to modify your selections if that was not your intent.

The act of exporting modifies the package version ID. The version ID does not change if you are only modifying the package contents.

Importing Packages

Generally, you perform package imports while ESM is running, because you want the behavior of the system to be affected immediately after a package import. Most package imports are small, incremental, and are of short duration. However, there can be large package imports,

which are later followed by periodic incremental changes to the packages. Large package imports can take up to 45 minutes to an hour.

Importing packages means you are sending resources from one ESM instance (the source) to another ESM instance (the destination). After the resources are installed in the destination, the import time will be the resources' new creation time. See also the related discussion on imported trends and their start schedules in the document, *ESM Best Practices: Trends*.

Best Practices for Importing Packages

During the package import, the important features of ESM must be available. Usually, import transactions succeed, but some could fail and result in roll backs of large imports. Rollbacks are automatic and are generally successful, and the system is returned to the state that it was in before the large import was attempted.

To safeguard your system against failures during large package imports:

- Perform a full database backup before a large package import.
- Have high processing power and large heap sizes if you are going to stress the system to a very high level.
- If possible, perform large package imports when the system is less loaded.
- If the package import fails, you can re-execute the import command and it should succeed.
- If you are importing a large package, you should *not* perform other tasks on the Console or the ArcSight Command Center while the import is going on. If necessary, ask another administrator to log in and perform those other tasks, while you are waiting for the package import to complete successfully.

- Maximum size for import packages:

By default, the maximum import size is 50,000,000 bytes. If the import package exceeds this size, the import will fail. You can change the size through this server property:

`xmlrpc.request.size.limit`

Increase the size in the `server.properties` file. See “Editing Properties Files” in the [ESM Administrator's Guide](#) for details on editing the `server.properties` file.

If the import introduces a new hierarchy for resources:

By default, if the import introduces a new hierarchy for resources, existing resources will keep their old hierarchies. This means you may have duplicate resources. This behavior is being controlled by the property called

`esm.manager.disable.resource.move`

By default, it is set to `true`.

If you are operating under very high loads, this default behavior can help prevent failure of import for large packages.

If you want your resources to move according to the new hierarchy, add the property to the `server.properties` file and set it to `false`.

See “Editing Properties Files” in the [ESM Administrator's Guide](#) for details on editing the `server.properties` file.

Where: Navigator > Packages tab

To import packages:

1. Click the  icon to import a package.
2. Navigate to the location of ESM packages. Choose an `.arb` file to import and click **Open**.
Importing the package copies it into the system where its package resource information is compared to any existing package with the same resource ID.
 - If the version IDs match, the import is aborted because the system assumes the import package is not different than the one in the system.
If you changed the one in the system, export it to give it a new version ID.
 - If the version IDs do not match, it continues to the next step.
3. By unchecking the **Install** box to the right of each package, you can choose to import a package without installing it. The default during import is to install. If you choose not to install an imported package, its resources do not appear. You can always install it later. See ["Installing or Uninstalling Packages" on page 742](#) for details.
4. Review the Import dialog for any conflicts. Each conflict displays one or more resolution options. To resolve a conflict, choose the preferred resolution option and click the **OK** button next to the options window. For more about resolving conflicts, see ["Resolving Package Conflicts" on page 745](#).
5. Click **OK** to continue. When the import is done, a Summary Report is displayed describing the packages that were imported.

Importing Packages Created by Other Users

Packages, like other resources, are always displayed under the user folder in which they were created. Upon import, the Summary Report shows the URI or full path into which the package was imported (for example, Packages Imported: `/All Packages/Personal/SomeUser's Packages/VPN Logins Reporting`). The import location is not configurable.

- If you log in with a different user name and import a package, you may or may not have write access to the package (depending on permissions).
- If you import the package with a different user name on a Manager that does not include an account for the package originator, you cannot see the imported package.
- If you recreate an account on the Manager with the same user name as the package originator, the imported package reappears.

Backing Up and Restoring with Packages

Although the package resource is the mechanism used to distribute ESM security use cases and Solution CIPs, packages are also designed as a backup mechanism for resources on running systems.

Resources can be part of more than one package. Therefore there are some behaviors associated with packages that may seem counter-intuitive and bear consideration.

ID Checking During Import

When a package is imported, there are some automatic ID checks:

1. The system looks for any other existing packages in the system with the same resource ID. This is the ID the system gave to the package when you created it.
 - If there are none, it imports the package and the process ends, unless there is a package in the same group with the same name.
 - If there is another package with the same resource ID, the evaluation goes to step 2.
2. The system compares the package version IDs for the importing and existing packages that share the same resource ID. The version ID is the ID the system gave to the package when the package was exported.
 - If the package version ID being imported matches the package version ID currently installed on the system, the package import process stops, because the system assumes that this package is already imported.
 - If the version IDs do not match, the evaluation goes to step 3.
3. The system checks each resource within the package to see if each version ID matches an installed resource with the same resource ID or URI (path and resource name).
 - If they match, the matching resource is not imported and the system checks the next resource.
 - If the version IDs do not match, the existing resource is over-written with the one in

the package being imported, unless you choose to import the package without installing.

If you import the package without installing it, the new package resource information is saved in ESM as an update and the icon changes to a small white up arrow in the lower left corner ().

Package Modifications

A package archive is a system data structure that contains the information defining a package and its resources. As you change a package and its resources, this file is not updated until you export the package. This enables the package to support the **Compare Archive with Current Package Contents** feature (from the package's right-click menu). This command allows you to see the packaged contents, for both the package last exported (the archive), and the current contents. The "Change Since Archive" column shows whether a given resource has been deleted, removed or modified.

When you export a package, the package's version ID is regenerated, regardless of whether the package attributes or any of its resources actually changed. This is not the case with the included resources; their individual version IDs only change upon export if the resource itself changed.

When you create a package, there is no version ID until you export it. Whenever you see a package with no version ID, that means there is no exported backup.

List Data

Active lists and session lists have two different uses as part of a package, and these affect how you would export the lists for backup purposes:

- The other resources in the package use the list to store data. The data is generated and used at run time. If, when you export the package, you do not need to save the data that accumulated in the list from the last run, use the **Export** package format.
- The list contains data that other resources in the package, such as rules, need specifically. If, when you export the package you must save the content of these lists, use the **Default** package format.

If you have a package containing some lists used only as containers and others with specific necessary data, use the Default format. The container lists would import again with data you do not need, but it is better than losing data you *do* need. Alternately, you could put the lists with required data into a second package using the Default format, and make this new package required for the first package, which uses the Export format.

Backup and Restore Summary

The version ID changes affect the results when importing a package in an effort to restore an existing package to a previous state.

- The system does not import a package backup if the version IDs of the package resource match.
- To make the existing system have a different version ID than the backup, you must export it again (being sure not to overwrite the backup you want to restore).
- If the existing package is bad and you have a good backup, you can always delete the existing package and then import the backup package.
- If you create a package that includes a top-level resource group (All Actors, for example), so you can back up the entire group, export the package often enough that all the recent changes are captured. If you ever delete such a package, it deletes the top-level group and every resource that has been added to the group, regardless of whether those new resources were added to the package. Be careful. Consider using the Children Only option.
- Generated version IDs do not identify which version ID is newer. For packages with the same resource ID, the system can only tell whether their version IDs are the same or different.
- To revert back to the last version imported, you must either delete the existing package or export it to some other location. Doing so sets a new version ID and places the export so as not to overwrite your backup. The backup package can now overwrite the existing package when you import it.
- Changing the name of an .arb file does not change the name, resource ID, or version ID of the package.
- If you currently have a package with version 1.1, and you want to import the backup package with version 1.0, there may be conflicts or other issues. See ["Resolving Package Conflicts" on page 745](#).

Installing or Uninstalling Packages

If you leave the Installed checkbox unchecked when you create a package, it is uninstalled. Uninstalled packages are not shown in the Normal view of the package tree. If you choose not to install a package when it is imported, and there was no other package with the same resource ID, the uninstalled package is essentially the same as a new uninstalled package.

However, if you imported a package for which there was already one in the system with the same resource ID, but a different version ID, and you chose not to install it, the system has two

different data sets for it. The existing package could be installed or uninstalled, but its Update Available attribute is now checked, which is indicated by an icon with a small white up arrow in the lower left corner (.

Where: Navigator > Packages tab

To install packages:



Caution: Packages can contain thousands of resources, therefore, installation may become a long-running transaction, up to 50 minutes or longer. While package installation is running, do not use the same administrator login to access and perform other administrative tasks on the or Command Center. Other administrator logins may access the or Command Center at this time, however, avoid running package installations concurrently with other administrators.

1. Locate the uninstalled package with the grayed out icon ().
2. Right-click the uninstalled package that you would like to install and choose **Install Package**.
3. Review the dialog for any conflicts. Each conflict displays one or more resolution options. To resolve a conflict, choose the preferred resolution option and click the **OK** button next to the options window. For more information, see ["Resolving Package Conflicts" on page 745](#).

Wait for the package installation to complete.

To uninstall packages:



Tip: While uninstalling a package, you might encounter a message that the package has a number of locked resources, and therefore cannot be uninstalled (or deleted, if you are deleting the package). Resources can be explicitly locked by their creators. Locked resources can also be system resources that are shared with other resources, and therefore cannot be uninstalled. There are two options:

- Unlock the user-created resources after verifying that the resources should be uninstalled with the package.
- If the locked resource should not be uninstalled, choose to skip the shared resources and proceed with uninstalling the package.

In the future, you may decide to exclude locked resources and shared system resources from any package to avoid this conflict.

1. Right-click the package to be uninstalled. The icon for installed packages is: .
2. Select **Uninstall Package**.

This command is disabled if the package is already uninstalled or if it is locked.

Uninstalling a package changes its icon appearance as grayed out, but the package remains in the system and can be installed again.

The resources that were in the package are removed from the system, even if they are also in another package. However, the details of these deleted resources are retained within the uninstalled package. When you reinstall the package, these resources are restored. Furthermore, they are also restored to whatever other packages they were in when you uninstalled this package.

You cannot add resources to an uninstalled package.

Deleting Packages



Caution: Before you delete packages, read the following information:

- Deleting a package that contains resources that maintain state—active lists with values, session lists, or trends—deletes the state information as well.
- When resources within the package are deleted, they are deleted even if they are contained in another package. Furthermore, if you delete a package that is required in another package (its icon has a red mark in the upper-right corner: ). The dependent package is useless without it, and is also deleted, along with all the resources in that dependent package.
- Before deleting a package, make sure the package excludes system resources, otherwise, these resources are deleted unless they are locked or they belong to a locked group. This can cause some serious problems especially when the system resources are Zones. If any system resources were deleted because the package in which they were included was deleted, re-import the package, edit the package so that system resources are excluded, then delete the package again.
- If there is even the slightest chance that you might need a deleted package or any of its resources at some point in the future, export it before deleting it.
- If you want to delete a package but not necessarily all the resources within it, remove the resources you want to save before you delete the package.

Where: Navigator > Packages tab

To delete a package:

1. Right-click the package to be deleted and select **Delete Package..**
1. Confirm that you want to delete the specified package.
2. If the package is still installed, you have the option to **Remove Resources in Package** or **Leave Resources**. If you leave resources, only the package itself is deleted. The resources that it contained remain in the resource tree. If you remove resources, all resources that the package contained are deleted from the system resource tree.

If the package is uninstalled, its resources have already been removed from the system, but deleting the package means you can no longer restore the removed resources.

See ["Backing Up and Restoring with Packages" on page 740](#).

Removing Resources from Packages

1. Click the **Packages** tab in the Navigator panel.
2. Right-click the package to be edited and select **Edit Package**. The Package Editor opens in the Inspect/Edit panel.
3. Click the **Resources** tab in the Package Editor.
4. In the upper half of the Resources tab, select the resource you want to remove. (A gray highlight on the entire row indicates the resource is selected.)
5. Click **Remove**.

Resolving Package Conflicts

Package conflicts can occur during install, uninstall, deletion, or importation of packages. Most package conflicts are resolved internally by the ArcSight Manager without the need for user intervention. However, some package conflicts prompt the administrator for an appropriate action from a list of options. This section describes two of these scenarios as examples.

If the ArcSight Manager detects package conflicts for a pending package **uninstall**, the Console provides choices for resolving the conflict and proceeding, or aborting the uninstall operation. The options provided depend on the type of conflict detected.

For example, if you attempt to uninstall a package that changed since it was installed, the conflict is indicated and you are asked to choose from the following Uninstall Resolution Options.

Option	Description
Create a new archive for package	Creates a new archive for the modified package (and retains the original).
Create new archive for remaining changed packages	Creates new archives for all changed packages before uninstall (retains all originals).
Continue without saving changes	Uninstalls this package without saving changes.
Uninstalls this and remaining packages without saving changes	Uninstalls all selected packages without saving changes.
Abort	Abandons the uninstall process and keeps the packages as is.

If the ArcSight Manager detects package conflicts for a pending package import or install, the Console provides choices for resolving the conflict and proceeding, or aborting the import operation. The options provided depend on the type of conflict detected.

For example, if you attempt to import a package with content that is older than the currently imported package, the conflict is indicated and you are asked to choose from the following Import Resolution Options:

Option	Description
Leave newer packages	Leaves the newer packages installed.
Never override newer packages	Completes the import but imports only packages that are newer than those currently installed.
Update packages	Imports the selected packages, and prompts for package conflict resolutions on a per-package basis.
Always update packages	Imports the selected packages, and overwrites newer packages if they exist.
Abort	Abandons the uninstall process and keeps the packages as is.

See the [ESM Administrator's Guide](#) for more information on working with the archive command and other utilities.

Chapter 27: Using Threat Detector

The Threat Detector feature of the ArcSight Console is activated when you get a license for the Threat Detector solution package. Threat Detector enables you to discover previously unknown patterns, which might pose a threat, and view them for analysis.

Threat Detector requires a separate license. Check your license agreement before using this feature.

Topics include:

- ["Threat Detector Overview" below](#)
- ["Threat Detector Life Cycle" on page 750](#)
- ["Creating or Editing a Profile" on page 750](#)

Threat Detector Overview

When finding threats by matching events against rules, you have to know the threat characteristics and create a rule that matches them. ArcSight Threat Detector enables you to search for threat patterns with known characteristics as well, but you can also find unknown patterns, where the only characteristic you specify is that the transactions are related and repeat.

The purpose of Threat Detector is to:

- Effectively search streams of potentially millions of events for patterns, which are simply repeating sequences of related events.
- Establish a baseline of patterns that represent normal event traffic and filter them out.
- Analyze what remains for threats.

In this way you can discover and investigate patterns that might represent new threats or threats whose characteristics are not known to you.

What Pattern Detection Provides

Threat Detector can automatically detect subtle, specialized, or long-term patterns that might otherwise go undiscovered in the flow of events. You can use Threat Detector to:

- **Detect day-zero attacks:** Threat Detector profiles are general enough that they can discover patterns that have never been seen before.
- **Detect low-and-slow attacks:** Low-and-slow attacks involve fewer events over a longer period. Profiles with longer time periods can capture these patterns.

- **Automatically discover patterns:** You can transform patterns into a rule set that is unique to your environment and more accurate than generic predefined rules.
- **Discover normal patterns:** New patterns discovered from current network traffic are like signatures for a particular subset of network traffic. You can specify which patterns are normal so that matching patterns can be eliminated as a threat.
- **Save a history of threat patterns:** Threat Detector can use event patterns that originate from or target an asset to categorize those assets. For example, a pattern of events from a machine that has an unauthorized program initiating a connection to an attacker (a back door) can be shown as a cluster. If you see this pattern originating from a new asset, it is a strong indication that the new asset also has a back door installed.

Use Threat Detector for preventive maintenance and early detection in your security management operations. Using periodic, scheduled analysis, you can continuously scan for new patterns over varying time intervals to stay ahead of new exploits.

Pattern Components

Events in a pattern share one or more common field values. For example, they could share the same source and target IP addresses, ports, host names, or other data.

The Threat Detector algorithm examines event components and identifies groups of related components as transactions. Discovered patterns list the components involved and the transactions containing common components. This data is output as a pattern resource. Components can relate to one another in several ways:

- **Related by session:** Session refers to a unique pair of source and target addresses. The events for which this pair are the same are in the same session.
- **Together in a sub-stream:** The event stream can be divided into sub-streams using a “group by” operation on a subset of event fields. This step can also take time of occurrence into account.
- **Together in time:** All the components occur together in a small time window.

Event components with some kind of relationship are grouped together as transactions, which then become potential candidates for patterns. The Threat Detector algorithm processes all the transactions it finds and produces patterns, depending on whether they satisfy one or more conditions that make them discernible as patterns.

Event components are subdivided into transactions in two major ways: time-based division, and event field-based division. These two methods can be combined.

Time-based division is based on timing constraints, and is very similar to the constraints used in defining rules. For example, the system creates a transaction at every division of an event stream. The event stream can be divided depending on the rate of occurrence of events and

changes in those rates. This works well for dividing event streams that display events in bursts of activity.


Event field-based division is very similar to doing a “group by” operation on event fields. Every related group of events is a sub-stream of the original stream of events. For example:

- **Based on source, target address, and port:** Suppose there are three distinct source addresses in the event stream. After doing a “group by,” three sub-streams are generated, each one originating from and corresponding to a unique source address.
- **Based on source and target address:** In this case, all the events that have the same source and target address belong to the same sub-stream.

How Threat Detector Works

Once the event stream is divided into transactions, Threat Detector identifies and groups events that occur together in multiple transactions. These events are sub-grouped by support level, which is the number of times that event occurred with its related events. A higher support number means that a pattern has occurred more frequently than others.

For example, consider the separate grocery purchase transactions, below. Several patterns emerge: Bread, butter, and jam were purchased together; as were milk and cereal. An analyst can draw conclusions from those patterns: these shoppers intend to make toast, or have cereal. Bread and strawberry jam also appear in two patterns and are a sub-pattern.



You can mask patterns you consider normal traffic so the system recognizes them and does not reevaluate them. For potential threat patterns that you want to watch for, you can build a rule

based on the pattern characteristics. When the pattern occurs, the rule triggers an action, such as notifying a group of users or running a command script.

Threat Detector Life Cycle

The creation and use of Threat Detector consists of three phases:

- Create a profile (see ["Creating or Editing a Profile" below](#))
- Generate snapshots (see ["Taking a Snapshot" on page 758](#))
- Investigate patterns (see ["Investigating Patterns" on page 765](#)).

Creating or Editing a Profile

A profile is a set of filters that define what fields to include in your pattern search, and the scope and properties of a pattern. It also specifies the time period to search. Profiles can be general or specific. Typically you would use several different profiles to define the parameters of snapshots, which collect all the events in the specified time frame and evaluates them according to the filters set in the profile.

To create a new profile:

1. In the Navigator panel, go to **Threat Detector** and click the **Profiles** tab.
2. Expand the **Profiles** resource tree.
3. If you are creating a profile, right-click a group in the resource tree and select **New Profile**.
If you are editing a profile, choose the profile you want to modify.
4. In the Inspect/Edit panel on the Profile Editor **Attributes** tab, you can modify most of the values (described below) and click **Apply**. Some values, such as version ID, are set by ArcSight and are not editable.

Profile Attributes

Property	Usage
Summary	A profile summary appears below the Attributes tab. The underlined items are values entered in the fields below.
Profile	
Name	Enter a descriptive name for your profile
Minimum Pattern Length	Type or use the up/down arrows to select the minimum number of unique associated events necessary to qualify the events as a pattern. The default value is 2 events.

Profile Attributes, continued

Property	Usage
Minimum Pattern Occurrences	Type or use the up/down arrows to select the minimum number of times for an event-association of the specified length to reoccur in order to qualify as a pattern. The default value is 2 occurrences.
Start Time	Select a time stamp expression for the snapshot start time. Expressions are described below. • \$Now The current time in the format hh:mm:ss. • \$Now - 1h The current time minus 60 minutes. • \$Now - 1d The current time minus 24 hours. • \$Now - 1w The current time minus 7 days. • \$Today The start of the current day (12:00:00). • \$Today - 1d The start of the current day at midnight (12:00:00) minus 24 hours. In other words, the start of yesterday. • \$CurrentWeek The start of the current week (Sunday 12:00:00). • \$CurrentMonth The start of the current month (the 1st 12:00:00). The format of start time is \$Now-<time>. The time is in increments of hours, days, weeks, or months.
End Time	Use the \$Now drop-down menu to select a timestamp expression for the snapshot end time. The formats are the same as for Start Time, above.
Events	

Profile Attributes, continued

Property	Usage
Event Fields, Source, Target	You can select one or more of these (event field, source, and target) for the pattern portion snapshot to display. Click in the data entry area and then click drop-down menu to see the field's chooser. In the Available Fields area, click the tab from which you want to choose. you can select one or more: - Field Sets - Local variables you created for this profile (see "Creating Local Variables" on page 756). - Fields and global variables that are relevant to a Threat Detector profile. In the Selected Fields section: - Use the up and down arrows to specify the order in which they appear. - Use the green alias icon to create an alias version. - Use the red X icon to remove one from the list. - You cannot specify date/time fields. - If you are going to add fields to a list, those fields must appear in this section (except the End Time field, which does not have to be here).
Restrict by Filter	Click the All Events drop-down menu to choose a filter from the Filters resource tree. The filter restricts the pool of events from which the snapshot is constructed.
Advanced	The check boxes in this section instruct the snapshot to capture elements pertaining to time, which can lend vital insight to a pattern. Tip: If you want to improve query performance and you do not need these options, leave them unchecked.
Record Time Order	This advanced option includes the time sequence of the events contained in patterns. For example, for a three-event pattern, it could record that A-B-C occurred 40 percent of the time, B-A-C 35 percent, and A-C-B 25 percent. Because event sequences can reveal intent, you can detect and act upon certain kinds of activity even sooner.
Split on Inactivity	This advanced option detects potentially meaningful decreases in activity between duplicate source/target pairs. It creates a break if there is a pause or significant drop in the number of times a particular pattern occurs. This treats occurrences of the pattern on either side of the break as separate instances. On analysis, a split on occurrences of the same source/target pairs means that there was a slow-down or break in occurrences. This enables you to discover patterns that happen repeatedly for one source/target pair.
Discovery Results	

Profile Attributes, continued

Property	Usage
Snapshot Retention Time	Click the drop-down menu to select how long you want the system to save a snapshot and its series of events. Snapshots retain all the needed components of the events and make them available during analysis. For example, when you drill down in an event and select “Show related events,” the events saved within the time frame set here will be searched for matches. The default retention time is 7 days.
Snapshot Group	Choose a group in the Snapshot resource tree in which to store the resulting snapshots. By default, the system adds the snapshot to the same folder you right clicked to add the profile.
Pattern Group	Choose a group in the Patterns resource tree in which to store the resulting patterns. By default, the system adds the pattern to the same folder you right clicked to add the profile.
Common	
External ID	An identification string suitable for, and which can be referenced by, systems outside ArcSight. Common applications of External IDs include appropriate naming for Case and Asset resources that are tracked in common with defect reporting or vulnerability-management systems. Your ArcSight administrator can advise you on the correct values for this field, if applicable.
Alias	An identification string suitable for referencing resources within ArcSight. A given alias appear in place of the resource's name everywhere it may be seen. Your ArcSight administrator can advise you on the correct values for this field, if applicable.
Version ID	If this profile came in a package or if you have exported it to a package, this is the package's version ID.
Description	A text description of the profile.
Assign	
Owner	The user with responsibility for the profile.
Notification Groups	The user groups to notify concerning changes to a profile.

To copy and paste a profile to another folder, select the profile to copy. Go to **Edit | Copy (Paste)** or use Ctrl + C (V).

To use one of these profiles, see ["Taking a Snapshot" on page 758](#).

Specifying Actions

The **Actions** tab enables you to select a trigger, then specify the action to take when that trigger occurs.

To specify an action:

1. Open the profile in the profile editor (double click the profile in the Navigator panel).
2. In the Inspect/Edit panel, click the **Actions** tab.
3. Before you add an action, specify when to take the action (the trigger). Select one of the following trigger options:

Trigger Option	Description
On Pattern Discovered	This specifies that the action be taken the first time a new pattern appears. Choose this option for assigning new patterns to an analyst to investigate.
On Pattern Re-discovered	This specifies that the action will be taken if a new pattern is repeated. Choose this option for ongoing operations.

4. Click **Add** and select one of the following options:

Threat Detector Actions

Action Option	Description
Annotate Pattern	In the dialog box, enter the following values and click OK : • Select a Stage from the drop-down menu. • Assign a user from the drop-down menu.
Set Event Field	In the dialog box, enter the following values and click OK : • Select a Field Set (or domain field set you created) from the drop-down menu. • In the event fields grid, set values for the event fields you are interested in.
Send Notification	Specify a notification group in the Notification Group drop-down menu. • Click Ack Required if those notified should acknowledge that they received notification. • Write the message to send in the Message field.

Threat Detector Actions, continued

Action Option	Description
Execute Command	In the dialog box, enter the following values and click OK: • Select an operating system platform from the drop-down menu. • Enter the command string. Use correct syntax; the system does not validate command strings. • Enter required parameters. For example, the archive tool needs the manager name, administrator name, and password. Specifying them lets the system execute the command without user intervention. • In the Action Type drop-down menu, select one of the following: Automatically run on manager: Initiates the command with no user intervention. Run on Manager with Console confirmation: Displays a confirmation dialog box in the ArcSight Console for the designated user before the command is initiated. Run on connector(s): Sends the command to the connectors that report the events.
Execute Connector Command	Specify a command to be executed at the SmartConnector reporting the events, such as pause or stop/start event flow. Enter the following values and click OK: • In the Connector drop-down menu, select the SmartConnector to execute the command. When you select an connector, the command field is populated with the commands available for that connector. • In the Command field, select the command for the connector to execute. The command may contain required parameters.

Threat Detector Actions, continued

Action Option	Description
Export to External System	You can export the pattern to an external tracking system, if you configured it to operate with ESM. Click OK .
Active List	You can add (or remove) a pattern to an active list, where its event details are available to other correlation tools for reference. • To add a pattern to an active list, select Add to Active List. In the dialog box, select an active list from the drop-down menu and click OK. • To remove a pattern from an active list, select Remove from Active List. In the dialog box, select an active list from the drop-down menu and click OK. • You cannot add fields to an Active List if they are not present in the Events section of the Profile. • You cannot add any date/time-based fields to an Active List since data/time fields cannot be included in the Events section of the profile.
Session List	You can add a pattern to a session list, or terminate a session list based on a pattern, where its event details are available to other correlation tools for reference. • To add a pattern to a session list, select Add to Session List. In the dialog box, select a session list from the drop-down menu and click OK. • To terminate a session list, select Terminate Session List. In the dialog box, select a session list from the drop-down menu and click OK. • You cannot add fields to an Session List if they are not present in the Events section of the Profile. • You cannot add any date/time-based fields to an Session List (except EndTime) since data/time fields cannot be included in the Events section of the profile. The End time displayed in the Add to Session List action is the time the entries are added to the session list.

5. The action summary will be displayed in the Actions tab. To remove lines that are not used, click **Hide Empty Triggers**.

Creating Local Variables

Click the **Local Variables** tab to manage local variables for this profile. These variables are available from the drop-down menu on the **Attributes** tab for event fields, source, and target attributes associated with the pattern.

From this tab you can:

- Add a variable.



Note: Variable names must be unique across resources. Local variables cannot share names.

- Edit a variable.
- Remove a variable.
- Make a variable global, which means it is available to resources outside this profile. If you make a local variable global, it moves it from the **Local Variables** tab to the **Fields and Global Variables** tab in the drop-down menu on the **Attributes** tab for event fields, source, and target attributes associated with the pattern. .

For more information on using local and global variables, see [Variables](#).

Threat Detector supports the following variable return data types:

- Byte
- Double
- Enumeration
- Integer
- Long
- Resource ID
- String
- Address

Therefore, function variables that return an unsupported data type are not supported. For example, the following functions or function categories are not supported:

- Non-SQL-mode variables
- Variables that return a list, such as `ActorByAccountID.AccountID` and variables that operate on multi-mapped active lists or overlapping session lists
- Variables that return a boolean value, such as the `Category Model` function `hasRelationship`

Adding Notes

You can keep track of changes made to a profile using the Notes feature.

To add a note to the profile:

1. In the Profile Editor, click the **Notes** tab.
2. In the Notes field, enter a note and click **Save** to log it in the Table/List tabs.

3. You can view notes as a table or as a list by toggling between the Table and List tabs. Reorder the table view by clicking the column header.

Deleting a Profile

1. In the Navigator panel, go to **Threat Detector** and click the **Profiles** tab.
2. Right-click a profile in the resource tree and choose **Delete Profile**.



Caution: You can delete or modify a profile if it has patterns and snapshots derived from it. However if you delete it, the patterns and snapshots that are derived from it no longer work and are not removed. If you modify it, they may not work as expected. Delete such patterns and snapshots when deleting their profile.

3. Click **Delete** in the confirmation dialog box.

Taking a Snapshot

A snapshot is a record of qualifying events that occurred over a specified period of time and evaluated according to the snapshot profile. When the Threat Detector algorithm runs on the specified data set, it displays the result as a graphic, which you can use for investigation and analysis.

You can generate snapshots manually, or run them on a schedule. You are likely to generate snapshots more frequently during the early stage of implementation, when you are establishing a baseline of normal patterns. Each snapshot is stored in the Navigator panel in **Threat Detector** on the **Snapshots** tab.


You can also discover patterns directly from active channels. Right-click a channel in the Navigator panel and choose **Discover Patterns**.

To take a snapshot:

1. In the Navigator panel, go to **Threat Detector** and click the **Profiles** tab.
2. Right-click a profile in the resource tree and select **Take Snapshot**.
3. In the Viewer panel, the system processes the snapshot request and shows each process as the Threat Detector engine runs. For example:

	Pattern Discovery run scheduled. Done!
	Building snapshot from events. Building snapshot from events. Processed 92357 out of 179698 events.
	Saving snapshot.
	Extracting patterns from snapshot.

- When the process finishes, the system displays the snapshot in the Viewer panel. The views are linked; click a node in the snapshot view to see its details in the patterns view.



The screenshot shows the ArcSight console interface. The top menu bar includes File, Edit, Views, Tools, Window, and Help. Below the menu is a toolbar with various icons. The main window is divided into two panels: the Snapshot view and the Patterns view.

Snapshot view: This panel displays a flowchart of events. The first node is "MS-SQL Worm propagation attempt" with "Snort" as the device and "Support = 45". This leads to the second node, "MS-SQL version overflow attempt", also with "Snort" and "Support = 45". This leads to the third node, "Attack From Suspicious Source", with "ArcSight" as the device and "Support = 18".

Patterns view: This panel displays a table of patterns. The table has two columns: "Name" and "Device Product". The patterns listed are:

Name	Device Product
MS-SQL version ...	Snort
MS-SQL Worm pr...	Snort
Slammer worm	

Below the table, the text "Slammer worm (suspicious)" is displayed.

Labels with arrows point to the "Snapshot view" and "Pattern view" panels.



Tip: If the pattern is empty, no events passed the profile's filter restrictions during the specified period. Adjust these profile specifications and generate the snapshot again.

Analyzing Snapshots


Use these options to analyze and respond to the patterns you discover in snapshots.

Snapshot Analysis Tools

Option	Usage
Create Rule	Use the Rules Editor to create a rule from a detected pattern of events or a selected event-level in the pattern hierarchy.
Show Related Events	Open a new channel filtered with a <code>matchesPattern</code> operator that uses the whole pattern, or event-levels, as its argument.
Show Event Graph	Graph the complete pattern or a selected event-level in the pattern hierarchy, to analyze using the ArcSight Console's visualization tools.
Inspect Pattern	The Pattern Inspector shows details, and you can click the Actions button to apply the options described in this table.
Investigate	You can create an active channel, or add a filter to the editor, using (or not using) the name of the selected event item in the pattern.
Tools	Choose one of the network tools ArcSight provides to explore the origin of the selected event item.
Annotate Pattern	You can mark the pattern with a workflow collaboration Stage and Assign it to a user for filtering by Stages and Users resources.

Exploring a Snapshot

Here is an example of a snapshot:



The upper part of the Viewer panel presents the snapshot view, which shows a hierarchy of related event nodes.


The lower part of the Viewer panel is the patterns view, which shows blocks of events from the hierarchy that are most closely related. Each block of events represents one specific path through the pattern hierarchy.

The example shows two patterns and a demarcation point (between support = 45 and support= 18). The top two events are the SQL worm. The last event is generated by the system. Threat Detector classified 18 of 45 sources as suspicious. There are 27 sources that ran the slammer worm in the network, but they were not added to the suspicious list. This discovery enables you to investigate why all 27 systems were not caught by the other surveillance mechanisms in place on your network. Determining that will help you to tighten your network security.

The “support” value for each node is the number of times that event occurred with its related events. The higher the number, the higher the item appears in the hierarchy. For example, in

the diagram below, there are two points at which there are sharp differences in support from one item to the next. This shift is called a demarcation point, and indicates a sub-pattern in a longer sequence.

The demarcation points (encircled in the figure below) indicate attack stages, and sometimes variations of the same type of attack on different network systems. For example, the SQL worm propagation attempt makes up 1000 of the 1122 hostile attempts. The demarcation point in the center of the graphic shows that there are two variations: attack from suspicious source, and UDP Packet tcpdump. This can indicate how different systems process the same type of SQL worm attack.



Arranging Elements in Graphic View

Use the buttons across the top allow you to zoom in, zoom out, and arrange the elements in different formations to give you better visibility of the overall pattern.

Tools for Rearranging Graphic Elements

Button	Control	Description
	Fit Content	Sizes the graphic to the available display space.
	Zoom in/ Zoom Out	Increases or decreases the size of the displayed graphic.
	Zoom Selected	Zooms in on a selected portion of a graphic.
	Hierarchic Layout	Presents nodes in a vertically descending cascade, similar to a family tree. Hierarchic layouts are appropriate when viewing relationships with a common root.
	Organic Layout	Arranges nodes based on minimum edge length, which tends to cluster items with a common relation. Clusters with items in common also tend to group together.
	Circular Layout	Hub-and-spoke arrangements with each node radiating edges to, or receiving edges from, the items with which it interacts. Circular layouts are most useful when multiple roots are present or there are a number of source-target relationships to clarify. If an organic layout is difficult to read because the edges are too dense, try a circular layout.
	Orthogonal Layout	Arranges items on the basis of logical connections, using electrical schematic-style right-angle layouts. These layouts are useful for clearly tracing connections and identifying node clusters.
	Overview	Opens a reduced rendering of the entire graph. You can drag the highlighted section in the reduction to move the displayed area in the main view.

In addition to the control buttons, you can drag items around in the Viewer panel while maintaining the connections. This can make the view clearer for overlapped items.


Scheduling a Snapshot

You can schedule a snapshot to be taken at intervals. The schedule frequency can be part of your daily analysis and operations. For example, as a best practice, you can run Threat Detector once a day to capture event patterns that happened over the last 24 hours. You can specify a longer period to find patterns with a longer term. To fully automate daily Threat Detector, add actions to a schedule, such as sending notifications, opening cases, or adding systems to an active list, if certain conditions are met.

Where: Navigator > Resources > Threat Detector > Profiles tab

Procedure:

1. Right-click a profile and select **Schedule Snapshots**.
2. On the Jobs tab, click **Add**.
3. In the Summary field at the bottom, select **Click here to set up schedule frequency**. This activates the Job Frequency dialog.



4. Click **OK** when you have set the frequency and time range.
5. Repeat as required to add more schedules for the same snapshot.
6. When you have added all the schedules for this snapshot, click **OK** at the bottom of the Jobs tab.
7. To add an action to be taken every time the profile is run, specify an action in the Actions tab of the profile editor, as described in ["Specifying Actions" on page 753](#).

Re-opening a Snapshot

If you have closed a snapshot in the Viewer panel, you can re-open it.

Where: Navigator > Resources > Threat Detector > Snapshots

1. Select the snapshot graph.
2. Right-click and select **Show Snapshot**.

When the snapshot's graphic has formed in the Viewer panel, you can click the icons at the top of the view to change its layout as described in ["Visualizing Resources" on page 714](#).

Deleting a Snapshot

Where: Navigator > Resources > Threat Detector > Snapshots tab

1. Right-click a snapshot in the resource tree and select **Delete Snapshot**.
2. Click **Yes** to confirm.

Investigating Patterns

When you take a snapshot, the Pattern view shown in the snapshot is also saved in the Patterns tab of the Threat Detector resource tree. You can use the Patterns tab to access more event investigation tools.

Investigating Patterns in the Snapshots View

Threat Detector gives you access to investigative tools from a series of buttons. These same tools are available from the right-click menu. The snapshot view and the patterns view offer most of the same investigative tools with a few specific differences. Right-click on any item in the graphical Snapshots view to open a new window within the snapshot view that contains details about the related events:

Right-Click Options for Pattern Investigation


Right-Click Option	Description
Show related events	Opens a new active channel in the Snapshots tab, filtered with a matchesPattern operator. This channel uses the pattern, or selected event-level in the pattern hierarchy, as its argument. To toggle back to the graphic view, click the Snapshot tab at the bottom of the snapshot Viewer panel.
Analyze in Channel	Creates a channel in a grid view that contains the associated events sorted according to Attacker Address, Name, and Target Address.

Right-Click Options for Pattern Investigation, continued

Right-Click Option	Description
Tools	Configure... includes the following options, and can be accessed directly through the larger Tools menu: • Nslookup - Resolves an IP address to a host name (domain name) and vice versa. • Ping - Determines whether a particular IP address is online and/or it tests and debugs a network by sending a packet and waiting for a response. • PortInfo - Lists standard usage such as WWW or FTP for a specified port number. • Traceroute - Shows the path from the ArcSight Console to the IP address selected in the grid view, reporting the IP addresses of all routers in between. • WebSearch - Search the Web through Google to find links to the keywords present in currently selected active channel grid view cells. • Whois - Looks up who is behind a given domain name; information might include addresses and telephone numbers. • Results - provides the results of running a network tool using the attributes of the selected pattern block For more information about network tools, see the online Help.
Create Rule...	Launches a Rules Editor in the Inspect/Edit panel. The rule you create here is stored in the Rules resource tree under the personal rules of the user who created it. For instructions about how to construct a rule, see "Creating Rules from Patterns" on page 770.
Show Event Graph	Displays the pattern as an event graph, which shows pattern components and their relationships in graphic form. For more information about ESM event graphs, see the online Help.
Show	Allows you to reset the graphic view with the following options: • Show all nodes - Displays the entire snapshot graphic. This is helpful if you have drilled down and wish to re-display the original snapshot. • Show all nodes containing selected items - Displays only the event hierarchy that contains the selected item. • Hide all nodes containing selected items - Displays all the event hierarchies that do not contain the selected item.

The following example in shows our sample pattern displayed as an event graph. To save space, the event graph consolidates items that have many members. In this case, the sample on the left shows the source address nodes consolidated into a single cluster with a single line representing the connections to each of the event name nodes.

To see the details and number of these connections, as shown on the right, uncluster the node by right-clicking the node and selecting **Uncluster** selected nodes.



Toggle between multiple views in the Snapshot window using tabs. Unclustering the source address nodes allows you to see the details of those nodes.

When you use the right-click menu to open a new view, it displays in a new tab within the snapshot pane. Use the tabs at the bottom of the pane to toggle between the views.

To close tabs in the snapshot view, right-click the tab at the bottom and select Close.

To rearrange open tabs in snapshot view:

1. Use the down arrow () to tile the open tabs horizontally, vertically, or to fit.
2. To select different views on an event graph, use the  button. For details about viewing event graphs, see the online Help.

Investigating Patterns in the Patterns View

You can re-open just the patterns view part of the snapshot in the Viewer panel.

1. In the Navigator panel, go to **Threat Detector** and click the Patterns tab.
2. Select one or more patterns in the resource tree, right-click the selections and choose View Pattern. This opens the Pattern pane in the Viewer panel.
3. You can take the same actions on the Pattern view as described in ["Investigating Patterns in the Snapshots View" on page 765](#).

In the Patterns view, you can click the Actions button or right-click a pattern, where you have the following options:

Toolbar Buttons in Patterns View

Button	Right-Click Option	Description
	Inspect Pattern	Opens the Pattern Inspector in the Inspect/Edit panel. For more information, see "Inspecting Patterns" on the next page .
	Create rule from Pattern	Launches a Rules Editor in the Inspect/Edit panel. The rule you create here is stored in the Rules resource tree under the personal rules of the creating user. For instructions about how to construct a rule, see "Creating Rules from Patterns" on page 770 .
	Annotate Pattern	Click this to open the Annotations dialog box. This allows you to escalate a pattern to another user for further investigation. For more information about how to annotate a pattern, see "Annotating Patterns" on page 772 .
Show ▾	Event Graph	Displays the events as an event graph, which shows interactions between two or more devices. For more information about how to use ESM event graphs, see "Graphing Attacks" on page 287 .
Show ▾	Related Events	Click this to open a grid view of the events contained in the Threat Detector snapshot.
Investigate ▾	Create Channel	Creates a channel based on the selected pattern block.
Investigate ▾	Add Condition to Editor	Enables you to edit the condition statements associated with this pattern block.

Viewing Patterns with Filter

You can view patterns assigned to a particular user or stage using Annotations (described in ["Annotating Patterns" on page 772](#)).

Where: Navigator > Resources > Threat Detector > Patterns tab

1. Right click a pattern and select **View Patterns with Filter**.
2. Use one or both of the following parameters for your search:
 - To filter for patterns assigned to a user, use the Select a User drop-down menu.
 - To filter for patterns assigned to a workflow stage, use the Select a Stage drop-down menu.

Inspecting Patterns

The Pattern Inspector provides you one more level of investigative control. If you decide that a pattern requires more investigation, you can use the Pattern Inspector to edit its details to be more descriptive for other users.

For example, you can rename the pattern from the default date and time of the snapshot to something more specific, such as “Potential worm attack.” Then you can add a description of the pattern so that another user can verify your findings.

Where: Navigator > Resources > Threat Detector > Patterns tab

To launch the Pattern Inspector:


Right-click a pattern in the resource tree and select **Inspect Pattern**.

Details of the pattern are displayed in the Inspect/Edit panel. Use the following sections as described below to tailor the pattern for further investigation:

Pattern Details

Section	Description
Summary	Use this section to modify the name of the pattern from the default date-and-time name to a more descriptive name. You can also add a description of the pattern to aid other analysts. The Profile field is not editable.
Items	Use the Recon drop-down button or right-click an item name to display the associated event details in a channel in the Viewer panel.
Snapshot	Use this drop-down menu to open patterns generated from the same profile definition so you can compare them.
Transactions	This table shows the source and destination data defined in the profile (address, port, host name, and so on) for the events involved in the pattern.
Time Spread	This table is only present if you selected Record Time Order in the profile. This table shows the details about the time spans involved between pattern occurrences. • Average - the average time between events in this pattern • Deviation - the difference in time spread between multiple occurrences of this pattern • Min - the minimum time between events in this pattern • Max - the maximum time between events in this pattern

The following Pattern Inspector shows item details and source/target transactions. You can rename a pattern to something more specific than the default date and time, and you can include a description.



Creating Rules from Patterns

You can create rules based on discovered patterns. Going back to our example, if Threat Detector finds a pattern between an MS-SQL worm propagation attempt reported by Snort, an MSSQL version overflow attempt, and an attack from a suspicious source, this indicates dangerous worm activity, and can create a rule to notify users or quarantine a server whenever the system detects traffic that matches this pattern. For additional information on creating and managing rules, see ["Managing Rule Actions" on page 532](#).


You can create rules from patterns in the Snapshot view in the Viewer panel, or in the Pattern Inspector in the Inspect/Edit panel.

- To access the Rules Editor from the Snapshot view:
Right click on any item in the hierarchy graphic and select Create Rule...
- To access the Rules Editor from the Snapshot Patterns view:
Right click on any item in the pattern block and select Create Rule.... You can also click the create rule button () in the button menu.

- To access the Rules Editor from the Pattern Inspector:
In the button menu, click the create rule button.

The Rules Editor opens in the Inspect/Edit panel showing the Attributes tab. Once the Rules Editor is open, do the following:

1. Follow instructions in ["Creating or Editing Rules" on page 509](#). You can also assign an external ID, alias, description, Version ID, owner, notification groups for the filter, and mark a resource as deprecated. Click **Apply**.
2. In the Rules Editor on the **Conditions** tab, the pattern's elements already appear in the common conditions editor. Modify the logic to express additional conditions for the rule to evaluate. For information, see ["Specifying Rule Conditions" on page 513](#).



Note: The OR conditions are intentional. OR is a more memory-efficient way to process rules than AND because it also applies a threshold value (the number of items involved) and distinct item names to track the components of the rule, rather than a blanket (join) approach.

3. At the **Aggregation** tab, set the number of matches and time frame for the rule.

4. At the **Actions** tab, set the actions for the rule to trigger when the thresholds are met.
 - a. Click **Hide Empty Triggers** in the top row. This reduces the list of available thresholds to those that are active (applicable to the conditions set in the rule).
 - b. Select a threshold from the list and click **Add**. Choose an action from the list that appears. See ["Rule Actions Reference" on page 539](#).
5. At the **Variables** tab, enter variables. Variables break down compound data fields into smaller parts so they can be sorted and acted upon. For example, you can break the 7-part timestamp field or a multi-value URI into component parts, which can be re-assembled in a more human-readable order, or sorted by component. For more about dependent variables, see the online Help and search for Variables.
6. You can keep track of changes made to a profile using the Notes feature:
 - a. In the Inspect/Edit panel, click the **Notes** tab.
 - b. In the Notes field, enter a note and click **Save**. The entry is logged in the Table/List tabs.
 - c. You can view notes as a table or as a list by toggling between the Table and List tabs. You can re-order the table view by clicking the column header.

Annotating Patterns

Annotation is a light-weight method to escalate a pattern to other users through your workflow system for analysis or investigation. You can use annotations instead of cases to escalate only one pattern. Use cases to escalate multiple patterns or if you use a third-party incident management system.

You can annotate patterns from the snapshot and Pattern views in the Viewer panel, or within the Pattern Inspector in the Inspect/Edit panel.

To access the Annotation Editor from the Snapshot Patterns view:

1. In the Navigator panel, go to **Threat Detector** and click the **Snapshots** tab.
2. Double-click the snapshot to display it in the Viewer panel.
3. Expand the pane so you can see the Patterns view at the bottom.
4. Right click any item in the pattern block and select **Annotate Pattern**. You can also click the Annotate Pattern button () in the button menu.

To access the Annotation Editor from the Pattern Inspector:

1. In the Navigator panel, go to **Threat Detector** and click the **Patterns** tab.
2. Navigate to the pattern and double-click it.

3. In the Inspect/Edit pane on the Pattern Inspector tab button menu, click the **Annotate Pattern** button.
4. In the Resource Annotation editor, enter the following values and click **OK**.

Field	Value
Stage	Select a stage from the drop-down menu. The default is Queued.
Assign to	Select a user from the drop-down menu.
Comments	Enter any comments to communicate to other ArcSight users.

Deleting a Pattern

1. In the Navigator panel, go to **Threat Detector** and click the **Patterns** tab.
2. Select one or more patterns.
3. Right-click the selected patterns in the resource tree and choose **Delete Pattern**.
4. Click **Yes** to confirm.

Threat Detector Usage Guidelines

Establishing a Baseline of Normal Patterns

Use broader profiles and more frequent snapshots to capture an example of all the patterns that occur as part of normal business practices. Identifying normal patterns takes time and investigation, and requires that you be familiar with traffic in your enterprise.

Once you have identified normal patterns, use annotation for moving them out of the analysis workflow. You can also use filters, but it is more reliable to move patterns by annotating them to a stage, such as Closed, because it assures that the pattern has been inspected and classified. For instructions about how to use event annotation to manage Threat Detector workflow, see ["Annotating Patterns" on the previous page](#).

Using Threat Detector in Routine Operations

Once normal patterns are identified and annotated so they are removed from the routine traffic flow, you can focus on the new patterns that are not yet classified. Routine operations consist of the following tasks:

- **Workflow.** As Threat Detector turns up new or unclassified patterns, a designated user needs to review them and start them through the workflow using the ESM annotations

feature. You can also schedule Threat Detector to run at intervals.

- **Investigation and analysis.** Once assigned to an analyst, the analyst can use the full array of ArcSight's investigation and analysis tools, including snapshot and pattern graphics, event graphs, filters, and rules, to determine the level of threat represented by the pattern.

During this investigation, it may be useful to drill down to the native device information to help identify the significance of a pattern. For example, if an event in a pattern was generated by Snort, you can retrieve the Snort rule number and look for its detailed explanation to obtain important event details.

- **Take action.** When a threat level is determined, the analyst can take a number of actions, such as use the ArcSight rule builder to take a prescribed action on this pattern and others that match it that may occur in the future; assign it to another user for follow-up; or close the pattern if it is deemed benign.

Performance Considerations

Threat Detector jobs can be resource intensive. Under high EPS, for example, greater than 15K, Threat Detector jobs can cause a degradation in performance, and may fail to return a matching result set. ArcSight recommends that you reduce the scope or frequency of Threat Detector jobs when running a system with high EPS.

Adjusting Threat Detector Memory

By default, Threat Detector limits its memory usage to about 4 GB of memory. However, if the search for patterns involves too many transactions and events, the task can run out of memory and abort. If the Threat Detector task aborts, a message to that effect appears in the ArcSight Console. Run the Threat Detector task again after increasing the Threat Detector memory usage limit. You can control the memory usage limit indirectly by changing the maximum number of transactions and events that can be held in memory.

For information, see "Adjusting Pattern Discovery Memory" in the Configuration chapter of the [ESM Administrator's Guide](#).

Chapter 28: Actors

The actors feature creates a real-time user model that maps humans or agents to activity in applications and on the network, making it possible to identify the actors behind events.

After the actor model is in place, you can construct category models to visualize relationships among actors and use those relationships for correlation.

This topic describes how to use the *actors* resources to model users and associate them with events. It also describes how to construct *category models* to depict relationships among actors.

Configuring Actors

Configuring actors requires a one-time setup procedure and minimal future maintenance if authentication systems are added, modified, or removed from your network. This setup procedure maps the user authentication systems you use in your network environment and the account IDs for each user on those systems.


1. **Install the Actor Model Import connector appropriate for your IDM.** For complete instructions about how to install the connector, see the relevant SmartConnector installation and configuration guide, such as the *SmartConnector Configuration Guide for Microsoft Active Directory Actor Model*. After installation, the connector polls the IDM and imports the user data into the Actor model.
2. **Identify the authenticators in your environment.** In preparation for configuring the authenticator mapping table, open the dashboard for automatically identifying the user authentication data stores running in your environment and their type:

/All Dashboards/ArcSight Administration/ESM/Configuration
Changes/Actors/Actor Administration

This dashboard is populated by the following query viewer, which looks for events with a value in the Authenticator field: /All Query Viewers/ArcSight

Administration/ESM/Configuration Changes/Actor/Actor Authenticators

The example below shows the value of the Attributes field for an active directory system configured as Active Directory:<domain>.com. Use this exact value, including punctuation, spaces, and capitalization, to populate the account authenticators mapping table described in the next step.




3. **Configure the Authenticators mapping table.** Using the information gathered in step 2, fill out the account authenticators mapping table provided at /All Active Lists/ArcSight System/Actor Data Support/Account Authenticators. The data you enter here must exactly match the values displayed in the Actor Administration dashboard.
 - a. In the Navigator panel, go to **Lists > Active Lists**. Right-click the active list /All Active Lists/ArcSight System/Actor Data Support/Account Authenticators and select **Show Entries**.
 - b. In the Account Authenticator Details tab in the Viewer screen, click the add icon ().
 - c. For each account authenticator data store, enter the following data:

Column	Description
Device Vendor	The vendor that supplies the authentication data store, such as Microsoft.
Device Product	Provide the application name of the authentication system, such as Active Directory.
Agent Address	The IP address of the reporting SmartConnector.
Agent Zone Resource	The zone in which the reporting SmartConnector resides.
Authenticator	Enter the exact value(s) returned for Authenticator in the Actor Administration dashboard from the previous step, including punctuation, capitalization, and spaces. Using the example shown in the previous step, the value you would enter in this column would be: Active Directory: arcsight.com

4. Leave Creation Time and Last Modified Time blank. ESM populates the Creation Time after you save the entry, and populates Last Modified Time as needed.

When you are finished, the Account Authenticators table should look something like this:



The screenshot shows a window titled 'Account Authenticators Details'. It contains a table with the following data:

Device Vendor	Device Product	Agent Address	Agent Zone Resource	Agent Zone Re...	Agent Zone Resource ID	Agent...
Microsoft	Microsoft Windows	192.0.2.0	<Resource URI="/All Zon...	2124	MX8HU5fsAABCCV7v-GN...	/All Z...
Microsoft	Exchange Server	192.0.2.1	<Resource URI="/All Zon...	2124	MX8HU5fsAABCCV7v-GN...	/All Z...
SAP	Security Audit Log	192.0.2.3	<Resource URI="/All Zon...	2124	MX8HU5fsAABCCV7v-GN...	/All Z...

Permissions Required to Use Actor-Related Data

By default, users in the Administrators group have full read/write access to the Actors feature and the other resources that actors depend on. The administrator can grant permissions for actors and the other resources upon which the actors feature depends to other users.

Permissions to create actors, actor channels, and category models:

- Read and write on /All Actors
- Read and write on /All Session Lists/ArcSight System/Actor Data and /All Session Lists/ArcSight System/Actor Data Support
- Read on /All Field Sets/ArcSight System/Actor Field Sets/Actor Base
- Read on the filters used to define the event ACLS for that user group, for example, All Filters/ArcSight System/Core
- Read and write on the group in which the new resource is being created

Permissions to view actors and category models, and monitor actor channels:

- Read on /All actors
- Read on /All Session Lists/ArcSight System/Actor Data and /All Session Lists/ArcSight System/Actor Data Support
- Read on /All Field Sets/ArcSight System/Actor Field Sets/Actor Base

Permissions to use actor global variables provided in standard content rules, active channels, and reports that leverage actor data:

Read access on the following resources and groups:

- /All Fields/ArcSight System/Actor Variables (either directly, or inherited from /All Fields/ArcSight System)
- /All Actors
- /All Session Lists/ArcSight System

- /All Active Lists/ArcSight System/Actor Data Support (for the authenticator active list)
- The appropriate group that gives all the queries used by a query viewer that leverages actor data
- The appropriate group that contains a query viewer that leverages actor data
- The appropriate group(s) for the filters used by any queries and query viewers that leverage actor data

In addition to these permissions on the actor-related resources themselves, read permissions are needed for any resources (such as filters, user-created actor global variables, and so on) upon which these actor-related resources rely.

For details about how to assign permissions to user groups, see ["Managing Permissions" on page 107](#).



Note: Best practice

Log out and log back in again for permission changes to take effect.

As a best practice whenever an administrator changes another user's permissions, the other user should log out and log back in again. This ensures that the new permissions are registered with the Manager, and the user can see the changes.

Viewing Actors on the

To navigate to Actors:

In the Navigator panel, select **Actors**. Here you will find the actors resource and the category models you can use to organize and visualize them.

In a typical workflow, actors are created automatically by installing the ArcSight Actor Model Import connector and configuring it to your IDM system. New actors added to the IDM are automatically created and existing ones updated with changes made on the IDM with every connection made between the Manager and the Actor Model Import connector as described in ["About the Actor Model Import Connector" on page 818](#).

For testing purposes, you can also create an actor individually using resources, or edit an existing one. For more about creating actors individually for testing purposes, see ["Creating and Editing Actors for Testing Purposes" on page 790](#).




Note: Console-created actors or those edited individually using Console resources do not update the user information stored in the IDM.

Communication from your IDM to the Manager is one way. Any actors that you add or existing ones that you update using tools are not added to the IDM system. Any changes you want to persist to the IDM should be made at the IDM, and the new actor information will be automatically imported into the actor model at the next Actor Model Import connector connection.

To view actor models:

You can view actor models with fewer than 1,000 members from the Navigator panel. Upon connection, the Actor Model Import connector creates the destination group in which the actors are placed based on the value set at the Actor Model Import connector. The example below shows three actors in a group called World-Wide Operations.



Viewing an Actor in the Actor Editor

To view the details of a particular actor in the Actor editor, double-click the actor, or right-click the actor and select **Edit**. Use the scroll bar to see all the actor attributes.

Inspect/Edit Actor: James Knight

Attributes | Notes

Actor

* U U I D	90758
* Full Name	James Knight
First Name	James
Last Name	Knight
Middle Initial	
IDM Identifier	Active Directory
D N	CN=James Knight, OU=Sales, DC=co...
Employee Type	full time exempt
Status	active
Title	Manager
Company	MyCompany
Org	SSD

(Name)
(Description)

Account Attributes + -

Authenticator	Account ID
Active Directory: companyname.com	CN=James Knight, OU=Sales, DC=co...
Active Directory: companyname.com	jknight@companyname.com
Active Directory: companyname.com	james_knight

Role Attributes + -

Role Name	Resource Name	Role Type
Manager	companyname.com	Business
Administrator	security software	Backup

OK Cancel Apply Help

Viewing Actor Account Attributes

The Account_Attributes table displays the unique account IDs attributed to this user by the various user authentication data stores relevant to this user. Like the base actor attributes, the values in this table are populated by values from the Actor Model Import connector for your IDM system.

Field	Description
Authenticator	The friendly name for the user authentication data store containing the actor's account ID
Account ID	The account ID for the actor

Viewing Actor Role Attributes

The Role_Attributes table displays role name, resource type, and role type for each role represented by the actor. The values in this table are also populated by values from the Actor Model Import connector for your IDM system.

Field	Description
Role Name	The name of the role or group, such as Administrators or Software Developer
Resource Name	The name of the resource in which the role is assigned, such as the active directory domain, identity management system, or application
Role Type	The role's category, such as Global Security Group, Business Role, or IT Role




Note: Because an actor can have multiple roles, a query viewer will display each role in separate entries. For example, if an actor has 4 accounts and 10 roles, running a query on this actor's accounts and roles will result in the Cartesian product of the accounts and roles: a total of 40 entries.

For information about queries, see ["Queries" on page 319](#). For information about query viewers, see ["Query Viewers" on page 340](#).

Viewing Actors in an Actor Channel

The actor channel is an active channel with a simplified header that displays the actor resources in your actor model.



The screenshot shows the 'Actor Channel' window. At the top, it says 'Active Channel:'. Below that, a filter is applied: 'Filter: Group ID = "World-Wide Operations"'. To the right, it says 'Total Actors: 12,456'. Below the filter, it says 'Inline Filter: No Filter'. A 'Radar' section is visible, which is currently empty. Below the radar, a table displays actor resources. The table has columns: Full Name, Title, Org, Department, Location, Employee Type, Status, Manager, and Email Address. The first row shows 'Jane Doe' as the Full Name, 'Account Representative' as the Title, and 'full time exempt' as the Employee Type, with 'active' as the Status.


Full Name	Title	Org	Department	Location	Employee Type	Status	Manager	Email Address
Jane Doe	Account Representative				full time exempt	active		



Note: The radar on an actor channel has no relevance. The radar applies only to event channels.

For actor models that contain thousands of members, actor channels present all the actors in your actor model in a single, scrollable view. You can apply local filters to actor channels to find actors with certain attributes.

If a group in your actor model contains more than 1,000 members, the actor tree in Navigator panel displays the message as shown:



You can also view actor models with fewer than 1,000 members in an actors channel.

To view an actor model in an actors channel:

1. In the Actors navigation panel, right-click an actors group and select **Show Actors**.

If your actor model contains more than 1,000 members, you can also double-click the message **“This group has more items than can be displayed. Double click to open a channel.”**



Note: When you select **Show Actors** on a group, the actor channel will only display the members of that immediate group. If the group has a sub-group, the actors in that sub-group will not be displayed in the actor channel.

To view the actors in a sub-group, right-click that group and select **Show Actors**.

2. In the Viewer panel, navigate to the actor channel.

The following sections describe the attributes of an actor channel, and how to interact with them.

Sorting Fields in Actor Channels

The fields shown on actor channels are from the Actor Information Field Set (/All Field Sets/ArcSight System/Actor Field Sets/Actor Information).

Sort fields in actor channels the same way you sort fields for event-based channels.



Note: You cannot sort multi-value columns, such as Account ID.

The names of sortable fields in column headers are indicated with a double-arrow icon . If a field is already sorted, an up  or down  arrow indicates the direction of the sort.

- To sort the list by a column, right-click over the column and select **Sort Column**.
- To reverse the sort order, select **Sort Column** again on an already-sorted column.
- To remove a sort, right-click over a sorted column and select **Remove Sort**.

For more about sorting columns in channels, see ["Sorting Events in the Active Channel" on page 238](#).

Actor Channel Options

There are several options available to take on actors from the tree view in the Navigator panel and from the grid view in the Viewer panel.

Actor Channel Right-Click Options from the Grid View

Option	Description
Export	Save the actor data in this actor channel as a CSV list.
Edit Actor	Open the selected actor to view its details in the event inspector.
Delete Actor	Delete the selected actor from the actor model. Caution: Make sure the actor is also deleted from the source IDM. Subsequent updates from the IDM that still contains this actor data can result in an unstable actor data set for this actor.
Add Actors to Category Model	Add the selected actors to an existing category model.
Add to Package	Add the selected actors to a new or existing package.
Report	Run a custom actor context report, or one using default values. For more information about actor context reports, see "Running Context Reports from an Actor Channel" on page 786 .
Find Actor in Navigator	Expand the containing group and highlights the selected actor in the Navigator panel.
Graph View	Display the actor in a resource graph in the Viewer panel.
Lock Actor	Locking is a common feature for all resources.

Filtering Actor Channels

There are two ways to filter the contents of an actor channel: adding a local filter to the resource itself, or applying an inline filter to one or more columns.

Adding a Local Filter to the Actor Channel Resource

You can add a local filter to the actor channel using the Active Channel: Actor Channel editor. This enables you to use the CCE to apply a filter locally to the selected actor channel. You cannot save a local filter that was added to an actor channel.



Note: If your filter uses a conditional variable function to display an actor list on a channel, the channel does not display any actor values. This is because conditional variable functions work for in-memory resources only. You can therefore use conditional variables for rules and data monitors, but not for queries and active channels.

For information about conditional variable functions, see ["Condition Functions" on page 1072](#).

To add a local filter to the resource:

1. Click the **Filter** link in the channel header.
This opens the Active Channel: Actor Channel editor in the Inspect/Edit panel.
2. In the **Attributes** tab, set the name and select the Actor field set you want to use. See the following guidelines:
 - a. **Name.** Replace the default name Actor Channel with a name that describes the channel, and perhaps the filter you want to apply to it, such as *Managers in World-Wide Operations*.
 - b. **Default Field Set.** By default, no field set is used. You can select a field set if you want to select a specific actor field in a particular field set. If you specify a field set, only the actor field sets are displayed, such as the Actor Base field set (Field Sets/Shared/All Field Sets/ArcSight System/Actor Field Sets/Actor Base). You can select this field set, or another actor field set created in your environment.
 - c. **Common Attributes.** Set any other common attributes you want for the actor channel. For a description of the data that goes in the Common section, see ["Common Resource Attribute Fields" on page 725](#).
3. In the **Filter** tab, construct the filter you want to apply. You can select any existing Actor field set, or apply a global variable.
For instructions about constructing a condition using the Common Conditions Editor (CCE), see ["Common Conditions Editor \(CCE\)" on page 887](#).
4. In the **Sort Fields** tab, select the columns by which you want the actor channel to sort. Fields that contain lists and multi-values cannot be sorted.
5. On the **Local Variables** tab, define any local variables you want to use to extract a particular value from a particular field. For instructions about how to use the Local Variables editor, see ["Variables" on page 1064](#).

6. Click **Apply** to apply changes to the actor channel displayed in the Viewer panel. Click **OK** to save the filtered actor channel.



Tip: Where to find saved actor channels

After you have modified and saved an actor channel, you can find it in the Active Channel area of the Navigator panel. Actor channels are saved with the suffix [Actor] after the active channel name, for example, *Managers in World-Wide Operations [Actor]*.

Creating an Inline Filter

Like event-based active channels, you can create an inline filter to operate on one or more columns to find actors with particular attributes in common.

For instructions about how to construct inline filters, see ["Filtering Active Channels with Inline Filters" on page 234](#).

In an actor channel, if you apply an inline filter to a specific column, the inline filter automatically becomes part of the actor channel's filter condition, as if you manually edited the actor channel and entered settings on the Filter tab. You have the option to save the actor channel with the new filter, or close the channel without saving the filter.

To save the filtered version of the channel, see ["Managing Actor Channels" below](#)

Managing Actor Channels

This topic shows how to save, edit, and view your saved actor channels.

To save an actor channel from the Viewer panel:

1. Right-click the active channel header and select **Save Active Channel As**.
2. In the Active Channels Selector, navigate to where in the Active Channels branch you want to save the actor channel and click **OK**.

You can also save an actor channel by opening the actor channel editor in the Inspect/Edit panel as described in ["Filtering Actor Channels" on page 783](#).

To edit a saved actor channel:

You can find saved actor channels in the Active Channel area of the Navigator panel. Actor channels are saved with the suffix [Actor] behind the active channel name.

1. In the Navigator panel, go to Active Channels.
2. Right-click the actor channel you want to edit and select **Edit Active Channel**.
3. Make modifications to the actor channel in the Active Channel: Actor Channel editor in the Inspect/Edit panel and click **OK**. For details about what to enter in the active channel editor, see ["Filtering Actor Channels" on page 783](#).



Tip: Where to find saved actor channels

After you have modified and saved an actor channel, you can find it in the Active Channel area of the Navigator panel. Actor channels are saved with the suffix [Actor] behind the active channel name, for example, *Managers in World-Wide Operations [Actor]*.

To view a saved actor channel:

1. In the Navigator panel, go to Active Channels.
2. Double-click the actor channel you want to view, or right-click it and select **View Active Channel**.

Investigating Actors

You can investigate events to identify the actor behind the activity represented in an event by running a context report from an event or actor channel. The actor context report looks at which actor is bound to the event you are investigating, and then run a report that will show activity for that actor.

Running Context Reports from an Actor Channel

From an actor channel, you can choose to run the report based on the following actor global variables:

- ActorByAccountID
- ActorByAttackerUsername
- ActorByCustomFields
- ActorByTargetUsername

The report will be populated if the actor global variable finds the values for the supported attributes, for example, account ID, custom field, attacker user name, and so forth. When the report is launched, it will use the actor global variable specified in the field set. If there is more than one actor global variable in the field set, the report will default to ActorByAccountID.



Note: Actor context reports will not show data if you are looking up actors using the ActorByUUID or ActorByDN global variable. These global variables are used only for internal actor lookups.


For context reports out of the actor channel, you have the following choices for running actor context reports:

- With default parameters:
 - Default time range: last hour
 - Default filter: correlation events only
- With custom parameters (you set these explicitly)
 - Start time
 - End time
 - Filter by

The following procedure shows the available options for running actor context reports from an actor channel.

To run an actor context report from an actor channel:

1. Display an actor channel and right-click an actor.
2. Select **Report** and then select one of the displayed report types.



If you choose a report type that ends in *with defaults*, for example, **Actor Context Report by Attacker Username with defaults**, the report is displayed with the following parameters:

- Default time range: last hour
- Default filter: correlation events only

If you choose a report type that does not end in *with defaults*, for example, **Actor Context Report by Attacker Username**, the following screen appears:

Name	Value
Common Parameters	
Report Format	html
Page Size	Letter [8.5x11 in]
Run as User	Select a User
Email to	
Email Format	Send URL
Email Subject	\$ReportName
Custom Parameters	
StartTime	\$Now - 1h
EndTime	\$Now
ActorResourceID	h11HHzysBABCAGnZTqyqk+g==
FilterBy	Correlation Events

☐ Save Output

OK Cancel Help

3. Set your custom parameters. For example, set StartTime, EndTime, and FilterBy. Keep the ActorResourceID parameter value; this is the value used to identify the actor of interest.
4. Click **OK**.

Investigating an Actor from an Event Channel

You can investigate an actor from an event channel in one of the following ways:


- By using the Show Actor option on an event that is related to an actor
This option is enabled if the channel contains ActorResourceID values, for example, ActorByAccountID.ID. Actor data is displayed on the Inspect/Edit panel.
- By running an actor context report on any active channel that has ActorResourceID values, for example, ActorByAccountID.ID.

Running an actor context report ("[Running Context Reports from an Actor Channel](#)" on [page 786](#)) provides additional options:

- Report with default parameters
- Report with custom parameters which you set explicitly

To show an actor related to an event:


1. Display an event channel.
2. Right-click an event and select **Show Actor**. If the channel does not use a field set containing an actor global variable, then the Show Actor option is disabled.



The edit panel displays details about the actor. See ["Viewing an Actor in the Actor Editor" on page 779](#) for an example of an actor edit panel.

To run an actor context report from an event channel:

1. Display an event channel.
2. Right-click an event and select **Report**.



3. Select the actor context report with default report parameters or the report that provides options to set report parameters for this report run.

The report is displayed on the 's Viewer panel.

Actor Context Reports in Standard Content

Actor context reports and supporting resources are available as part of the standard content. For locations and descriptions of the reports and related resources, refer to the [ArcSight Administration and ArcSight System Standard Content Guide](#).

If needed, you can modify the resources upon which these context reports are based. Refer to the following topics in this guide:

- For details about modifying reports, see ["Defining Report Attributes" on page 394](#).
- For details about working with queries, see ["Query Settings" on page 321](#).
- For details about working with report templates, see ["Using Report Templates" on page 426](#).

Creating and Editing Actors for Testing Purposes

For testing purposes, you can create an actor using resources, or edit an existing actor resource. If you are manually creating actors, manually enter data in the fields you are interested in tracking.

In a production environment, the Actor Model Import connector automatically populates the actor attributes the connector has been configured to send, based on values set at the source IDM. The IDM may not use or store data for every field. To learn more about the values the Actor Model Import connector can be configured to send, see the Actor Model Import connector documentation for your IDM system, for example, the *SmartConnector Configuration Guide for Microsoft Active Directory Actor Model*.

Important Points to Consider About Making Manual Changes to Actors

If you are creating, editing, or deleting actors that the IDM system sent through the Actor Model Import Connector, consider the following factors:

- Actors you create using the are not sent back to the IDM. The flow of data is one way from the IDM through the connector into ESM.
- Any changes you want to persist to the IDM should be made in the IDM itself. Any new actor information will be automatically imported at the next scheduled Actor Model Import connector connection.
- If you made manual changes in the to actors imported from the IDM, these changes will be overwritten the next time the Actor Model Import connector sends updated data for the same actors.
- If you manually deleted an actor attribute, that attribute will not be updated by a subsequent update from the Actor Model Import connector, unless the connector report includes an updated value for the attribute that you deleted.
- You should be careful about using the to delete actors sent by the IDM, especially if the actors still exist in the IDM, because it is possible that the actor will not be updated during a subsequent import.

Creating Actors for Testing Purposes

Before proceeding, review the information in ["Important Points to Consider About Making Manual Changes to Actors" above](#).

To enter basic Actor attributes:

1. In the Navigator panel, go to **Actors**. Right-click the **All Actors** group (or any group under **All Actors**) and select **New Actor** to launch the Actors editor.

You can also launch the Actors editor by going to **File > New > Actor**, or by clicking the New Resource icon () and selecting **Actor**. If you used the **File > New > Actor** menu option, the actor is added in the Unassigned folder. Later, you can move the unassigned actor to an existing group.

2. In the Actor Editor in the Inspect/Edit panel's Attributes tab, enter values for the required fields, UUID and Full Name. Enter any other relevant attributes. All attributes are treated as data type string. Use the scroll bar to see all the Actor attributes.

To enter values in the Account Attributes table:

In the Account Attributes table, add all the unique account IDs attributed to this user by the various authentication data relevant to this user.




Note: About values in the Account Attributes table

- In a production environment where the IDM sends data through an Actor Model Import connector, this table is automatically populated with the user account ID and authenticator values that the Actor Model Import connector is configured to send to ESM.
- For tips about tools to use to find user account IDs and authenticator information, see ["Configuring Actors" on page 775](#).
- In a test situation, or any situation where an actor has been added manually using tools, you manually populate the account attributes you are interested in tracking.

1. Click **Add** () to make the fields editable.
2. In the *Authenticator* column, enter an identifier for the user authentication data store, for example, *Active Directory: mycompany.com*. This is a friendly name that will help admins and other users identify which data store is the authentication source.
3. In the *Account ID* column, enter the user's account ID used in that authentication data store, for example, *john_doe*, *jdoe*, or *john.d*.

With each entry, the next set of fields becomes editable. Add as many data store authenticators and account IDs as are relevant. For example, an entry for an Active Directory authenticator could be *Active Directory: companyname.com*. Following is an example of a completed *Account_Attributes* table:



Authenticator	Account ID
Oracle	meerkat
MS Windows	jmeerkat
MS Exchange Server	jmeerkat@jaygroup.com

4. To remove an entry, click anywhere on the row you want to delete and click the Delete icon ().

To enter values in the Role Attributes table:




Note: About values in the RoleAttributes table

- In a production environment where the IDM sends data through an Actor Model Import connector, this table is automatically populated with the user role values the Actor Model Import connector is configured to send.
- In a test situation, or any situation where an actor has been added manually using tools, you manually populate the role attributes you are interested in tracking.

1. Click the Add icon () to make the fields editable.
2. In the *Role Name* column, enter a role name, such as Administrator, User, Approver, or Manager.
3. In the *Resource Name* column, enter the application to which this role applies, for example, SAP or Microsoft Exchange. In the *Role Type* column, enter what type of role it is, such as whether it's an IT role or a business role.

With each entry, the next set of fields becomes editable. Add as many user roles as are relevant. Following is an example of a completed Role_Attributes table:



Role Name	Resource Name	Role Type
DBA	Oracle database	IT
Administrator	MS Exchange Server	IT

4. To remove an entry, click anywhere on the row you want to delete and click the Delete icon ().

Editing Actors for Testing Purposes

This section contains instructions for editing the actors you have created manually for testing purposes. In a production environment, actor changes should be managed automatically through the Actor Model Import connector.

1. In the Navigator panel, go to Actors.
2. Double-click an actor (or right-click an actor and select **Edit Actor**) to open the Actor Editor in the Inspect/Edit panel.



Tip: How to find an actor among thousands of actors

If a group in your actor model contains more than 1,000 members, view the actors using an actors channel. In the Actors navigation panel, right-click the group and select **Show Actors**.

For more about creating and viewing actors in an actors channel, see "[Viewing Actors in an Actor Channel](#)" on page 781.

3. Refer to the topic "[Viewing an Actor in the Actor Editor](#)" on page 779 for details about what to enter in the editor's fields.

Deleting Actors

The actors feature is designed to reflect the latest state of your Active Directory data as sent through regular updates from the Actor Model Import connector. Changes made to actor data in your environment should be made first at the Active Directory; updates from the Actor Model Import connector then reflect these updates in ESM.

If you delete a user in Active Directory, the actor remains in the actor model, and its status is updated to `Deleted` in IDM. This gives you the opportunity to control the permanent deletion of actor data from ESM, or keep the information based on your requirements.

- If you created small actor groups manually for testing purposes, right-click the actor group or specific actor and select **Delete**.
- If you are deleting a large number of actors imported from Active Directory, this process can take an extended period of time. It can lock the Console and the deletion may not even be completed. Follow the instructions in the topic, "Best Practices for Deleting a Large Number of Actors" in the *Active Directory Model Import Connector Configuration Guide*.

Leveraging Actor Data Using Variables

You can create a local or global variable that focuses just on actor base and list fields. This enables you to make a specific value derived from actor data available for use in actor-related resources: actor field sets, actor queries (used both in reports, query viewers, and trends), and actor channels.



Note: Use of velocity expressions is not supported in actor fields and in local or global variables for actor fields.

Creating an Actor Global Variable

Variables derive particular values from existing data fields. The global variables feature enables you to define your variables only once, and then re-use it in multiple places wherever conditions can be expressed. Global variables work with the actors feature so you can build user correlations.

To create an actor global variable:

1. Launch the global variable editor: In the Navigator panel, go to **Field Sets**. On the Fields & Global Variables tab, right-click a group and select **New Global Variable**.
2. In the Attributes tab, give the global variable a name, and specify **Actor Global Variable** as the variable type.
For details about the fields in the Global Variable Editor Attributes tab, see ["Global Variable Editor: Attributes Tab" on page 612](#).
3. In the Parameters tab, specify the parameters you want to set for the actor global variable.
 - a. In the Function field, select a category, then a function appropriate for the data you want to extract from the actor fields.
 - b. In the Arguments section, select the fields or resources to which you want to apply the function. Enter the other relevant arguments for that function.
 - c. To test the result returned by the parameters you selected, enter test values and click **Calculate** to test the results of the actor global variable.
4. In the Local Variables tab, you can optionally add a local variable to the actor global variable, which will extract a value from a field that you want to use in the overall actor global variable.

For details about how to create a global variable using the global variable editor, see ["Creating or Editing a Global Variable" on page 610](#).

For details about the functions available to local and global variables, see ["Variable Definition Fields" on page 1066](#).

Creating an Actor-Based Variable in Another Resource

Actor-based variables are only applicable to actor-based resources. You can add a local variable based on an actor field to the following resources:

- Actor active channels
- Field sets

- Global variables
- Queries (available to reports, trends, and query viewers)

To create actor-based local variables:

1. In the resource editor **Local Variables** tab, click **Add**.
2. Provide the following information, and then click **OK**:
 - a. Enter a name for the local variable.



Note: Ensure that the name is unique across resources. Local variables cannot share names.

- b. Select a function that is compatible with the actor field whose values you want to leverage in the variable.
- c. In the **Arguments** section, select fields and add values relevant to the actor data you want to leverage.
- d. In the **Preview** section, enter test values and click **Calculate** to test the results of the actor global variable.

Creating and Using Category Models

After you have actor information created, you can make logical groupings to represent relationships among actors and actor attributes using category models.

Category models can reflect direct actor relationships, such as reporting hierarchies, or relationships between actors who share common attributes, such as actors in a particular location. For reporting hierarchies, your model can consist of a top-to-bottom structure (by Manager), or its reverse (by Assistant). Category models can also reflect relationships between actors using custom attributes defined by the user.

You can use category models to visualize these relationships, then leverage the data gathered in them using the `HasRelationship` function in local and global variables.

Memory Recommendations for Using Category Models

Category models can be resource intensive on run-time processing memory, depending on the size of your actor model and the nature of the relationships you are modeling. For best results, adjust Java Heap Memory Size in the `setup` script to at least 1 GB.

To adjust the Java Heap Memory Size on the :

1. If running, close the .
2. In the directory <ARCSIGHT_HOME>/bin/scripts/, make a backup of the startup script file:
 - Windows: console.bat
 - Unix: console.sh
3. Open the Console startup file (console.bat or console.sh) in a text editor, and change the default maximum heap size value from -Xmx512m to -Xms1024m.

For example (value to change is highlighted):

For Windows:

Change the line

```
set ARCSIGHT_JVM_OPTIONS=-Xms64m -Xmx512m -XX:MaxPermSize=84m -
```

to

```
set ARCSIGHT_JVM_OPTIONS=-Xms64m -Xmx1024m -XX:MaxPermSize=84m -
```

For Unix:

Change the line

```
ARCSIGHT_JVM_OPTIONS="-Xms32m -Xmx512m -XX:MaxPermSize=84m "
```

to

```
ARCSIGHT_JVM_OPTIONS="-Xms32m -Xmx1024m -XX:MaxPermSize=84m "
```

4. Save the updated Console startup file.
5. Restart the .

Creating Category Models

You can create three types of category models depending on the type of relationships you want to represent:

- **Actor-to-actor.** Actor-to actor category models establish direct or indirect relationships between actors themselves, such as reporting hierarchies. This category model is also called a dual-field category model.
- **Model by actor attributes.** Actor attribute category models are a way to group actors who share one base actor attribute in common, such as location, department, or country. This

category model is also called a single-field category model.

- **Model by user-defined attributes.** User-defined category models are a way to group actors who share one or more attributes that are outside of the schema, for example, users who come in on Saturdays, users who play racquetball, or users who take public transportation. This category model is also called a manually-created category model.



Caution: Manually-created category models will not be included in an export of the actor resource. For more information about exporting resources, see ["Creating or Editing Packages" on page 731](#).

To create a category model:

1. In the Navigator panel, go to Actors and click the **Category Models** tab.
2. Right-click an existing group and select **New Category Model**.
3. In the Category Model Editor's Inspect/Edit panel, name the category model, select its type, and select the fields by which you want it to model. For details about what fields to populate for the type of category model, see the following topics:
 - ["Creating Actor-to-Actor Category Models" below](#)
 - ["Creating Actor Attribute Category Models" on page 800](#)
 - ["Creating User-Defined Category Models" on page 802](#)
4. Depending on the type of category model you create, use the Data tab to view the members of the category model, or use the Attributes tab to define the attributes by which you want to model users.
5. Optional: To add information in the Notes tab, refer to ["Using Notes" on page 79](#).
6. Click **OK** to save the category model and close the editor, or click **Apply** to save the category model and leave the editor open.

Creating Actor-to-Actor Category Models

Actor-to-actor category models establish direct or indirect relationships between actors themselves, such as reporting hierarchies. The categorization is based on what data you want to track using the Parent Field, and how to look up the actors for populating the model through the Child Field.

When creating actor-to-actor category models, enter the following values in the Attributes tab of the Category Model editor:

Actor-to-Actor Category Model Attributes


Attribute Field	Description
Name	Enter a name for the category model. This name will appear in pick lists and wherever category models can be referenced in conditions. Spaces, underscores, and hyphens are allowed.
Create From	Use the Create From field to select the type of category model. For an actor-to-actor category model, select Actor Fields. After the category model is saved, this field becomes read-only.
Field Set	By default, the system uses the Actor Base field set, because only actor-based fields are relevant for category models. You can also select a user-defined actor field set. The field set selected here defines the fields available for the parent and child field choices and defines the columns available for the table below the graph view (for more about the graph view, see "Viewing Category Models in Graphs" on page 805). After the category model is saved, this field becomes read-only.
Relationship Name	Use this optional field to describe the relationship you want the category model to show. For example, if you want the category model to show managers and their direct reports, you could enter Direct Report to identify the relationship between the manager and the direct report. The value you enter here appears as a mouse-over tool tip on the relationship lines that connect the parent and child fields in the category model graph view.

Actor-to-Actor Category Model Attributes, continued

Attribute Field	Description
Parent Field	This field enables you to establish which Actor data field to use to build a hierarchy of relationships. From the drop-down menu, select Manager or Assistant as the parent field. For example, if you are building an actor-to-actor category model that shows top-down reporting relationships, select Manager to produce a category model of every manager identified in your IDM data. The resulting model displays managers at the top. A Parent Field of Assistant displays an inverted hierarchy with lower-level actors appearing at the top. Note: Only the Manager and Assistant fields are supported for the Parent Field when building actor-to-actor category models.
Child Field	In Child Field, the UUID (or DN, if used) is the unique identifier the system uses to look up the actors and populate members who are related to Parent Field. From the drop-down menu, select UUID or DN as the unique identifier to correctly determine who the members are in a particular structure. (DN is specific to the Active Directory IDM.) Note: Only the UUID and DN fields are supported in the Child Field when building actor-to-actor category models. For example, if you selected Manager in the Parent field, and the actor's Manager field is populated by a UUID value, then select UUID as the Child Field here. Likewise, if the actor's Manager field is populated by the DN value, then select DN here. The following example scenario explains how the category model is created based on the Parent Field and Child Field values. Example scenario: Assuming you are building an organizational chart with managers at the top node: • Actor A has a UUID = 1234. Actor A is the manager of Actor B and Actor C. • Actor B and Actor C's values for Manager = 1234, which corresponds to Actor A's UUID. • In building this category model, use Parent Field = Manager and Child Field = UUID. The Manager looks up Actor B and Actor C's Manager field, which has Actor A's UUID. It then creates Actor B and Actor C under Actor A in the resulting category model.
Delimiter	The delimiter field does not apply to the actor-to-actor category model.

For a description of the data that goes in the Common section, see ["Common Resource Attribute Fields" on page 725](#).

Use the Data tab to view the members of the group in tree form.



You can also view the group hierarchy in a resource graph. Right-click the category model and select **View Category Model**. For details, see ["Viewing Category Models in Graphs" on page 805](#).

Creating Actor Attribute Category Models

Actor attribute category models are a way to group actors who share one base actor attribute in common, such as location, country, or any actor attribute that can possibly have a hierarchical groupings.

When creating actor attribute category models, enter the following values in the Attributes tab of the Category Model editor:

Actor Category Model Attributes

Attribute Field	Description
Name	Enter a name for the category model. This name appears in pick lists and wherever category models can be referenced in conditions. Spaces, underscores, and hyphens are allowed.
Create From	Use the Create From field to select the type of category model. For an actor attribute category model, select Single Actor Field . After the category model is saved, this field becomes read-only.
Field Set	By default, the system uses the Actor Base field set, since only actor-based fields are relevant for category models. You can also select a user-defined actor field set. The field set selected here defines the fields available for the parent and child field choices, and also defines the columns available for the table below the graph view (for more about the graph view, see "Viewing Category Models in Graphs" on page 805). After the category model is saved, this field becomes read-only.
Relationship Name	Use this optional field to describe the relationship you want the category model to show. For example, if you want the category model to show employees by location, you could enter Location to identify the relationship between the actor and the group he is associated with. The value you enter here appears as a mouse-over tool tip on the relationship lines that connect the actor and the attribute they're being modeled by in the category model graph view.

Actor Category Model Attributes, continued

Attribute Field	Description
Parent Field	From the drop-down menu, select the attribute that you want to model the users by. For example, if you are building an actor attribute category model that categorizes all the actors by their location, select Location.
Child Field	The child field does not apply to single-actor field category models.
Delimiter	Enter the delimiter you used in the actors' Delimiter attribute. The default is the forward slash (/). The Delimiter is used to denote the hierarchy of values, from top to bottom, in the attribute you are tracking in Parent Field. Example scenario: • Actor A has Location = /USA • Actor B has Location = /USA/California/Mountain View • Actor C has Location = /USA/California/Mountain View The delimiter used to denote a hierarchy is /, therefore, in the category model editor, set Delimiter = /. The resulting resource graph produced by these values has three levels: USA, California, and Mountain View.



Tip: A combination of delimiters builds a hierarchy if one is found

If the attribute from which you are creating the category model contains multiple values with more than one type of delimiter, for example, a URL, such as

`http://www.greatcompany.com`

include all the delimiter characters in the Delimiter field. For example:

`://.`

This indicates that the dot (.) is the delimiter used to separate all the elements of the URL into the following hierarchy:

`http`


`www`

`greatcompany`

`com`

For a description of the data that goes in the Common section, see ["Common Resource Attribute Fields" on page 725](#).

Use the Data tab to view the members of the group in tree form.



Based on the hierarchical relationship in the example above, the value for the two actors' Location field is entered as /US/CA/MountainView.

You can also view the group hierarchy in a resource graph. Right-click the category model and select **View Category Model**. For details, see ["Viewing Category Models in Graphs" on page 805](#).

Creating User-Defined Category Models

User-defined (or manually-created) category models are a way to group actors who share one or more attributes that are outside of the actor schema, for example, users who come in on Saturdays, users who play racquetball, or users who take public transportation.

The user-defined groupings can be created in hierarchical fashion. For example, users to who take public transportation can be further classified into those who take the train, those who take the bus, and those who take a ferry. For user-defined category models, the hierarchy evaluation is based on the actor's UUID value.

1. When creating category models based on user-defined attributes, enter the following values in the Attributes tab of the Category Model editor:

Attribute Field	Description
Name	Enter a name for the category model. This name will appear in pick lists and wherever category models can be referenced in conditions. Spaces, underscores, and hyphens are allowed.
Create From	Select the type of category model from this field. For a user-defined attribute category model, select Manually . After the category model is saved, this field becomes read-only.

Attribute Field	Description
Field Set	By default, the system uses the Actor Base field set, since only actor-based fields are relevant for category models. You can also select a user-defined actor field set. The field set you select here defines the columns available for the table that appears below the graph view (for more about the graph view, see "Viewing Category Models in Graphs" on page 805). After the category model is saved, this field becomes read-only.
Relationship Name	Use this optional field to describe the relationship you want the category model to show. For example, if you want the category model to show actors who take different types of public transportation, you could enter <code>Commutes By</code> to identify the relationship between the actor and the group he is associated with. The value you enter here appears as a mouse-over tool tip on the relationship lines that connect the actor and the attribute they're being modeled by in the category model graph view.

Parent Field, Child Field, and Delimiter don't apply to this category model.

- For a description of what to enter in the Common fields, see ["Common Resource Attribute Fields" on page 725](#).
- Use the Data tab to define the attributes by which you want to group users, and to add actors to the category model.

For example, if you want to create a hierarchy of users that take different types of public transportation, do the following:

- In the Category Model editor at the Data tab, click **New Group**.
The name of the new group is automatically highlighted so you can give it a relevant name, for example, `Public Transportation Commuters`. Press the Enter key to save the new name.


To rename a group at any time, right-click the group and select **Rename**; or click the **Rename** button. After entering the new name, press Enter.
- Add actors to the category model group. You can add actors from the Navigator panel and from an actors channel in the Viewer panel.

From the Navigator panel: Drag and drop actors from the navigator panel into the category model group. You can drag and drop multiple actors at a time.

From an Actors channel in the Viewer panel: You can view any group of actors in an actors channel in the Viewer panel. An actors channel is the only way to view groups with 1,000 or more members. Select the actors you want to add to the category model group, right-click, and select **Add to Category Model**.

- To select multiple actors in a row, use **shift + click**.
- To select multiple actors out of sequence, use **ctrl + click**.

- c. To create a sub-group of the first group, such as Train Commuters and Bus Commuters, click **New Group** again. You can also right-click the existing group (or right-click anywhere in the editor panel and select **New Group**.
 - By default, the new group is made a child of the first group.
 - You can make the new group a parent group by dragging and dropping it to the desired location, or click **Move Out** (or right-click the group and select **Move Out**).
 - You can make a parent group the child of another by dragging and dropping the group to the desired location, or click **Move In** (or right-click the group and select **Move In**).



- You can have unlimited parent and child nodes.
4. View the category model in a resource graph. Right-click the category model and select **View Category Model**. For more about viewing category models as graphs, see ["Viewing Category Models in Graphs" on the next page](#).

Managing Category Models

To edit a category model:

1. In the Navigator panel, double-click the category model to open the Category Model editor. Or right-click the category model and select **Edit Category Model**.
2. Make edits to the category model attributes.
3. Click **OK** to save the category model and close the editor; or click **Apply** to save but leave the editor open.

For details about the category model fields for the different types of category models, see ["Creating and Using Category Models" on page 795](#).

To copy or move a category model:

You can move or copy a category model the same way you move or copy any resource.

1. In the category model resource tree, navigate to an asset and drag and drop it into another group.
2. Choose **Move** to move the category model, **Copy** to make a separate copy of the category model, or **Link** to create a copy of the category model that is linked to the original category model.

If you choose **Copy**, you create a separate copy of the category model that will not be affected when the original category model is edited. If you choose **Link**, you create a copy of the category model that is linked to the original asset. Therefore, if you edit a linked category model, whether the original or the copy, all links are edited as well. When deleting linked category models, you can either delete the selected category model or all linked category model copies.

To delete a category model:

1. In the Navigator panel, right-click the category model and select **Delete Category Model**.
2. At the confirmation dialog box, click **Delete**.


Viewing Category Models in Graphs

To fully visualize a category model, view it as a graph. Category model graphs are very similar to resource graphs, only instead of modeling relationships among all resources, category model graphs render relationships among attributes of actor resources.

To view the category model graph:

In the Navigator panel, right-click a category model and select **View Category Model**.

The Viewer panel displays the graph. By default, all top-level nodes are displayed in collapsed form. To expand the nodes, click the plus (+) sign.



Working with category model graphs:

The category model graphs are displayed on the Viewer panel. As with all resource graphs, There is a set of command buttons at the top of the view and a parallel set of commands available by right-clicking the graph itself.

The table below shows the Category Model Toolbar Buttons and Right-click Commands.

Category Model Graph Options

Command	Button	Description
Fit Content		Size the model to the available display space.
Zoom In / Zoom Out		Increase or decrease the size of the displayed model. The system is optimized to show the entire category model graph in the available space of the viewer panel. If a category model has many nodes and members, its elements can appear very small. To zoom in, click the zoom in icon (+ magnifying glass). To zoom out, click the zoom out icon (- magnifying glass).
Zoom Selected		Zoom in on a selected portion of the model.
Hierarchic Layout		Present nodes in a vertically descending cascade, similar to a family tree. Hierarchic layouts are appropriate when viewing event relationships that have a common root.
Organic Layout		Display nodes in an arrangement based on minimum edge length, which tends to cluster nodes that relate to a common node. Likewise, node clusters with nodes in common will also tend to group together.
Circular Layout		Position nodes in hub-and-spoke arrangements with each node radiating edges to, or receiving edges from, the nodes with which it interacts. Circular layouts are most useful when multiple roots are present or there are a number of source-target relationships to clarify. If an organic layout is difficult to read because the edges are too dense, try a circular layout instead.
Orthogonal Layout		Arrange nodes on the basis of logical connections, using electrical schematic-style right-angle layouts. These layouts are very useful for clearly tracing connections and identifying node clusters.

Category Model Graph Options, continued

Command	Button	Description
Overview		Open a reduced rendering of the entire graph. You can drag the highlighted section in the reduction to move the displayed area in the main view.
Hierarchy Tree		Open a complete list of the nodes as seen on the category model editor's Data tab. Click a node in the list to scroll to that node in the main view.
Print		Print the displayed model.
Export to JPEG		Create and save a JPEG-format copy of the current image.
Add Graph View to Case		Add the current graph view to a case you select. Choosing this option opens the Case Selector dialog, where you can browse cases. Select a case to which to add the current graph view and click OK on the Case Selector dialog. The graph view is added to the selected case as an attachment, accessible on the Attachments tab in the case editor for that case.
Help		Display the relevant online Help topic.
Expand One Level/ Collapse One level		Expand <i>all</i> collapsed nodes to display nodes one level below. Collapse one level of all expanded nodes. This feature works only if any node was not manually expanded or collapsed previously.
Plus (+)/Minus (-)		Expand or collapse a single node on the graph.
Increase/ Decrease Node to Node Distance		Increase or decrease distances between nodes by small increments. This feature works on expanded nodes.
Single-person icon		Denotes an individual actor.
Two-person icon		Denotes a group of actors.

Every time a node is expanded or collapsed, the entire category model graph re-sizes to fit into the available space of the viewer panel



Note: Actors with no value for the field used to define the category model do not appear in the model

If an actor does not have a value for the field that was used to build the category model, that actor will not appear in the model.

For example, if you build the category model on the Office attribute, and if an actor does not have a value in the Office field, that actor is not represented on the category model view.

For more about resource graphs in general, see ["Visualizing Resources" on page 714](#).

Leveraging Category Model Data Using Variables

Use the HasRelationship function in local and global variables to leverage data represented in a category model.

Local and global variables are available in resources that use conditions: active channels, filters, rules, data monitors, and queries. Local variables are available for use only in the resource for which the variables are defined; global variables can be re-used in multiple condition-based resources.

To leverage data represented by a category model in a variable:

1. Launch the variable editor.
 - **Local variable:** From the channel, filter, rule, query, or data monitor editor, click the Local Variable tab. Click **Add** to launch the *Add Local Variable* dialog box.
 - **Global variable:** In the Navigator panel, go to Field Sets. On the Fields & Global Variables tab, right-click a group and select **New Global Variable**.

In the Attributes tab, specify **Event** as the variable type.

For details about the fields in the Global Variable Editor Attributes tab, see ["Global Variable Editor: Attributes Tab" on page 612](#).

2. Select the function category **Category Model** and the HasRelationship function. In the Arguments section, select the category model whose data you want to leverage and specify the parent and child field or group.

Arguments for HasRelationship Function	
Field	Description
Name	When creating local variable based on category model data, provide a friendly name for the variable. This name is used anywhere the variable is applied (CCE, resource field selectors).
Function	From the Function drop-down menu: - Select the category Category Model. - Select the function HasRelationship. - Click OK.
Category Model	Browse to and select the category model from which you want to leverage data.

Arguments for HasRelationship Function	
Field	Description
Parent Field or Group	Navigate to the field or single-value variable you want to use as the parent. Use the Field/Group drop-down to indicate whether the parent is a field (single attribute) or a group.
Child Field or Group	Navigate to the field or single-value variable you want to use as the child.
Inherit All Related Actors	Select true for the variable to consider all the actors in a related hierarchy. For example, VP > Director > Manager > direct report. Select false for the variable to consider only direct relationships between actors. For example, Manager > direct report.

For details about working with the Global Variable editor, see ["Creating or Editing a Global Variable" on page 610](#).

For details about local variables and the functions available to both local and global variables, see ["Variable Definition Fields" on page 1066](#).

Chapter 29: Reference Guide

The topics that follow provide information on resources, components, and terms used throughout the guide. Topics are organized alphabetically, and introduced and defined in a style meant to help you get more drill-down information about a term quickly and easily. Unlike a standard “glossary,” however, many of these topics present quite a bit of in-depth information including conceptual and reference material. These topics are cross-referenced (linked) extensively with the rest of the Help topics and vice versa.

Access Control Lists

ESM uses Access Control Lists (ACLs) to manage user group permissions. ACLs define which user groups have permissions to which resources, and to which components such as rules, reports, and filters. (See also ["Editing Access Control Lists \(ACLs\)" on page 107.](#))

User groups can have inspect (read) permissions, edit (write) permissions, or both. If a user group has inspect permissions, they can read the resource. For example, the users in the group can see the resource and related information through the . If the group has edit permissions, they can write to or change the resource, such as writing or editing a rule or report resource.

Resources, too, can have inspect (read) permissions, or edit (write) permissions. Resources, like user groups, are managed as groups and not as individual resources. Therefore, a resource can only be accessed if a user group has access to the resource's group. Permission to inspect or edit resources is granted when the user logs in, and the resource only appears in the if the logged in user has inspect permissions.



Note: Best practices:

- Log out and log back in again for permission changes to take effect.
- Whenever an administrator changes another user's permissions, the other user should log out and log back in again. This ensures that the new permissions are registered with the Manager, and the user can see the changes.

Resource ACLs

Resources have ACLs to help you manage user permissions based on the resource. You can use the resource ACL to determine which user group can access it. You can control which user group has access to inspect or to edit any resource, such as rules, cases, and reports. (See also ["Editing Access Control Lists \(ACLs\)" on page 107.](#))

Events are also available to user groups based on resource ACLs. For example, you can control which user group has access to a filter by adding the group to the filter's ACL and giving them

inspect or edit permissions. If you no longer want the group to have permissions to that filter, you can edit the group's permissions or remove the group from the filter ACL. In this example, the user group listed on the filter ACL with inspect permissions can see events from that filter in the . Those without permissions cannot see any events from that filter.



Note: By default, a custom user group inherits the ACL settings from the parent user group. If a user group has no access to any filters, it is as if the group's specified filter in the ACL editor were *Filters\Shared\All Filters\ArcSight System\Core\No Events*. Therefore, users from this group do not see any events in the following resources:

- Active channels using event filters
- Queries on events, and therefore query viewers and reports based on those queries

If users need to see event data, make sure their user group has access to the applicable filters. See ["Adding or Removing Enforced Filters" on page 113](#) for details.

Events are also extracted from the database based on ACLs. For example, when users generate reports, events extracted from the database are based on ACLs. Therefore, only data that users have access to is retrieved and all data may not be included in the report. Report ACLs can only provide events if the user generating the event has the permissions to view those events. For example, if user group A has permissions to view events from filter A and user group B does not, user group B cannot be able to extract event values from filter A when running a report; the report comes back empty. However, since user group A does have permissions to filter A, user group A's report comes back with the values from filter A.



Note: The Resource ACL display shows relationships between users and groups, and how permissions are acquired for each of the user groups. Child groups inherit permissions from parent groups.

For example, consider the following set of ACLs for assets.

Resource	R	W
/All Users/Administrators	☒	☒
/All Users/Default User Groups	☒	☐
/All Users/Default User Groups/Analyzer Administrators	☐	☒

In this scenario, the following permissions apply:

- A user logged in as Administrator (belonging to the group /All Users/Administrators) has read and write permissions by virtue of being in the Administrators group.
- All users have read permissions because they belong to the group /All Users/Default User Groups by default.
- A user logged in as an Analyzer Administrator has both read and write permissions because these are inherited read permissions from the parent group (/All Users/Default User Groups) and get write permissions per the Analyzer Administrators child group

Active Channels

Almost all event-related views are **active channels**. Also, several types of resources related to assets and cases are shown as active channels.

Rather than simply flowing events through as received, or capturing a fixed set of events for replay, a channel is in effect a live, on-going event query. Because it is continually re-evaluated, the set of events collected in a channel can continue to change (due to reporting latency), even when defined with a fixed time-bracket.


In other words, active channels are definitions for collections of events; definitions that are always freshly re-evaluated so the resulting sets are as valid as the data received up to that moment. Because the active channel continuously refreshes with live events, you should not use an active channel to track event counts because these counts vary across different timeframes. On the other hand, a replay channel's event count is based on that specific replay session (with the specific sessionID). Instead of active channels, consider using query viewers or reports to track event counts.

The queries that define active channels are composed, at a minimum, of time parameters; other filter conditions of the usual sorts can also apply. You find and use these queries in the Navigator panel's Active Channels resource tree. You create these definitions through the **File>New>Active Channel** command and can refine them using inline filters and the Active Channel Editor. Once defined and displayed, you can manipulate the order, format, and content of these views with all the familiar features of the Console.

Query viewers are provided as a quick alternative to active channels, better suited to some scenarios. See ["Query Viewers" on page 340](#) for more information.


Active Channel Views

Each individual view is one **rendering** of an active channel, whether it is a grid view or chart view. Individual views are represented by the tabs you see at the bottom edge of the Viewer panel. Channels are represented by the tabs at the top of the Viewer panel, that group together individual views.



Active Channel Headers

The channel name and statistics line appears at the top of active channel views. These statistics are event-severity indicators for the view. The indicators show the current number of events in the view for each of the priority categories. You can click these indicators to instantly filter the channel to show only the selected priority.



The **Filter** status line describes the filter conditions the channel is currently using.

The Radar display in active channel headers indicates the activity taking place in the channel, in graphics that represent units of time horizontally. The time automatically adjusts to accommodate the scope of the channel. The broader the scope, the smaller the graphical units become.

You can open and close the Radar display with the **Plus (+)** and **Minus (-)** button at the right end of the Filter line.

With keyboard actions, you can control the contents of a grid view using its Radar display.

Click, **Shift+click**, **Ctrl+click**, or drag to select one or more contiguous or non-contiguous bars in

the display. You can also drag selection borders left or right to adjust a span further. The grid then shows just the events the selection represents.

Comparisons

You may want to note that the Manager handles active channel traffic through the database. This means that the content is persisted, but may involve processing delays that cause an active channel to show information later than a more direct method such as the data monitors in dashboards. Conversely, data monitor traffic resides only in memory and is subject to loss or abbreviation by server restarts.

See the topic ["Monitoring Active Channels" on page 224](#) to learn more about active channel tasks. You may also want to compare active channels to active lists as analysis tools.

Active Channel Views for Assets and Cases

The shows assets, vulnerabilities, asset categories, scanner reports, and cases in **active channels**. You can leverage the power of channels for asset management, including use of filters, field sets, and better sorting capabilities.

About Actors

A critical factor in having situational awareness is knowing who is doing what with resources on your network, when they're doing it, and how. This awareness is critical for maintaining network security and demonstrating compliance with the increasing requirements of regulatory standards.

Identity management systems (IDMs) enable IT security professionals to protect their assets while granting different levels of access to a range of users, such as full-time employees, part-time employees, employees with certain security clearances, partners, contractors and customers.

However, following exactly what a specific person is doing across all the resources on your network can be difficult, because each user will have different account IDs and roles on different systems and applications. Examples of different information used to identify a given user include badge IDs (physical access devices), MAC addresses (for devices assigned to a specific person), email addresses, user names, Distinguished Names (particularly for Active Directory-related events), and so on.

The Actors feature maps humans and their activity to events from applications and network assets by leveraging user attributes defined within identity management systems and

correlating them with user account information from the user authentication systems on your network. Correlating user identifiers from the event traffic that reflects their activity throughout the day makes it possible to ensure that users are doing role-appropriate activity across the assets in your organization, and to detect and track inappropriate access and suspicious activity.

The Actors feature works in conjunction with the SmartConnector for Active Directory Actor Model Import which regularly polls your Active Directory. This system automatically maintains an up-to-date actor model you can use to correlate users and their roles with their activity on the network.



Note: ArcSight supports actor models with up to 500,000 actors

ArcSight supports actor models with up to 500,000 members. Supporting a large actor model can require special configuration. For details, see “Tuning Guide for Supporting Large Actor Models” in the “Configuration” chapter of the [ESM Administrator's Guide](#).

Also, see the `searchindex` command in the “ArcSight Commands” chapter of the [ESM Administrator's Guide](#). The `searchindex` command is a utility that creates or updates the search index for resources in the database.

If the actor model is in place, use the modeling and visualization tools (*category models*) to show direct and indirect relationships among actors in the actor model. You can use this model to group and visualize users in your organization in numerous ways, such as reporting structures, organizational units, or role-based functions, then use these relationships as parameters in user-defined monitoring, analysis, and correlation.

For testing purposes, you can also manually add actors. You can also import or redefine views of user groups and relationships with category models.

Actor channels and navigating thousands of actors

Actor channels present all the actors in your actor model in a single, scrollable view. Like active channels, apply local filters to actor channels to find actors with certain attributes.

Actor channels are the only way to see actor models that contain 1,000 or more members because display space in the Navigator panel is limited. You can also use actor channels for viewing actor models with fewer than 1,000 members.

For more about actor channels, see ["Viewing Actors in an Actor Channel" on page 781](#).

Viewing relationships among actors using category models

After you have configured actor information, you can make logical groupings to represent relationships among actors and actor attributes using category models.

Category models can reflect direct actor relationships, such as reporting hierarchies, or relationships between actors who share common attributes, such as actors in a particular location. Category models can also reflect relationships between actors using custom attributes defined by the user.

To visualize relationships using category models, leverage the data gathered in the models using the `HasRelationship` function in local and global variables.

Refer to these related topics:

- ["Creating and Using Category Models" on page 795.](#)
- ["Viewing Category Models in Graphs" on page 805.](#)
- ["Leveraging Category Model Data Using Variables" on page 808.](#)

Using actor global variables to identify actors from events

The actor data stored in the Actor Resource Framework coupled with actor global variables make it possible to identify an actor from any given event, then correlate that activity with other activity or attributes of that actor. The ability to identify an actor from a given event and correlate that activity with other events involving that actor and attributes of that actor, such as location and role, make it possible to verify that an actor's activity across the network is appropriate.

Standard content provides a series of actor global variables that are part of the Actor Resource Framework, which identifies and stores actor-related data from events in the look-up tables of the Actor Resource Framework. You can also use these global variables in your own correlation content. For more about using the Actor Resource Framework global variables, see the [ArcSight Administration and ArcSight System Standard Content Guide](#) that comes with this ESM release.

You can construct your own actor global variables. For an outline of this process, see ["Leveraging Actor Data Using Variables" on page 793.](#)

Using Standard Content to Track Actor Configuration Changes

Standard content also provides a set of coordinated resources that track actor configuration changes, such as when actors are created, updated, and deleted.


For more about this standard content, see the [ArcSight Administration and ArcSight System Standard Content Guide](#) that comes with this ESM release.

How the Actors Feature Works

ArcSight SmartConnectors normalize event data from hundreds of different devices on a network into a common data schema. The Actors feature normalizes user identity information stored in different formats in different authentication data stores to create a complete profile of data used to identify each user on your network in various contexts.

As shown in the following example, a model import connector imports data from an identity management system, such as Microsoft Active Directory. For a complete list of supported identity management systems, contact Customer Support.

In the following example, the actor data comes in from the Microsoft Active Directory system using a model import connector. Events arrive from applications that all use different data stores to authenticate user activity, which all use different account IDs to identify the user John Zed. The activity is identified as belonging to the same actor. That actor is represented as JOHN.



The actors feature works in conjunction with an Actor Model Import connector to import user data from an identity management system and to normalize user data in event traffic.

The actors feature is supported internally using the Actor Resource Framework, a series of internal look-up tables maintained by regular updates from the Actor Model Import connector.

As part of setting up the actors feature, you also configure an applications and *authenticators active list* to identify the mapping between the applications in your network environment and the data stores they use to authenticate users. In the example shown above, Windows Server Active Directory is the authentication data source for Microsoft Exchange and SAP Real-Time Security.

The following diagram shows a detailed look at how the Actors feature works. When events arrive at the Manager, resources that use conditions or select fields invoke one or more of the actor global variables provided in standard content. These global variables and the actor data maintained in the Actor Resource Framework provide several ways to identify actors using whatever user identity attributes are available in events arriving from different applications from across the network.

The global variables first look up the authenticator using the device-specific data, such as vendor and product information in the event, then look up the relevant user information from the Actor Resource Framework tables to positively identify the actor. For details about the Actor Global Variables in standard content, see the [ArcSight Administration and ArcSight System Standard Content Guide](#).

Resources leverage system-provided actor global variables to look up actor identity attributes maintained in the Account Authenticators table and the Actor Resource Framework.

About the Actor Model Import Connector

The ArcSight Actor Model Import connector support bulk import of user accounts from Microsoft Active Directory. Refer to the configuration guide for the *ArcSight Actor Model Import Connector for Microsoft Active Directory*.

The Actor Model Import connector imports the user data into the actors resource, where it is leveraged by the infrastructure that identifies and tracks user activity. Correlated and normalized data about user activity is then available for monitoring and investigation, further correlation, and reporting.

The actor model used to describe users is automatically populated with the attributes configured for it by the Actor Model Import connector when connecting to the connector.



Caution: Actor Model Import connector should be configured with all attributes you are interested in tracking before initial connection.

During Actor Model Import connector configuration, make sure that all the attributes you are interested in tracking are configured. After actor information is imported into ESM, and the Actor Model Import connector has an updated list, existing actors in ESM are not updated.

To update existing actors with new or removed attributes sent from the Actor Model Import connector after an actor model has already been imported, first delete the actor group, then re-import the actor data.

For details about how to delete an existing actor group, see ["Deleting Actors" on page 793](#).

The following table lists the attributes that the actor model supports. The Actor Model Import connector administrator configures the Actor Model Import connector with the attributes from this list that it will send to populate the actor model. Not all IDM systems support all these

attributes. An actor resource is only populated with the attributes configured by the Actor Model Import connector administrator.

Single-value attributes		Multi-value attributes
UUID	Manager	Account
First Name	Assistant	
Middle Initial	Email Address	- Account ID - Authenticator
Last Name	Location	Role
Full Name	Office	
IDM Identifier	Business PhoneMobile Phone	- Role Name - Resource Name - Role Type
DN	Fax	
Employee Type	Address	
Status	City	
Title	State	
Company	Zip Code	
Organization	Country Or Region	
Department		

In addition to the basic single-value attributes, each actor will likely have multi-value attributes, specifically multiple account IDs, and multiple roles, which are tracked using your IDM system. These multi-value attributes can appear differently in events coming from different devices. In some cases, such as a non-IT-related role, the information is not included in event data at all, but is still valuable information to help identify users and correlate their activity to help ensure appropriate behavior and access to resources hosted on the network.

Troubleshooting Errors with Actor Model Imports

It is possible that during the actor import process from the Actor Model Import connector, one or more actor import files containing data for multiple actors may not have imported successfully into the Manager. This might happen because of network connection problems, an out-of-memory error, or some other problem that caused the import of that file to fail.

In such cases, there is an archive file in `$ARCSIGHT_HOME/archive/webservices` for each actor import file that failed to import successfully. Each such archive file is created with the file extension `.bad`.

If an actor file did not import as expected, or during routine maintenance, check the `$ARCSIGHT_HOME/archive/webservices` directory for actor files that failed to import.

The `.bad` archive file contains all the missing actor information, and you can use the ArcSight Archive utility to import that file individually from a command line on the Manager system. For instructions about how to run the Archive utility to import an archive file, see the topic "ArcSight Commands" in the *ESM Administrator's Guide* and locate the description for the archive command.



Tip: Using the Archive utility:

- To see a list of commands available with the Archive utility, include -h (for “help”) in the archive utility command script.
- If the archive file name starts with a dash (-), rename the file before running the Archive utility to ensure that the command works.

For details about the Actor Model Import connector and how to configure it, see the Actor Model Import connector documentation for your IDM system, for example, the *SmartConnector Configuration Guide for Microsoft Active Directory Actor Model*.

Active Lists

You can use **active lists** to create a configurable data store that can hold information derived from events, or other sources.

Active lists can monitor activity based on any rule-driven combination of event attributes or set of custom fields. For example, active lists are very useful for tracking suspicious or hostile IP addresses as well as targets of attacks that may be compromised.

You can populate active lists manually when necessary (adding entries from grid views or the Active List Editor), or use active lists in conjunction with rules specifically tailored to work with them. Rules can dynamically add and remove entries on active lists, thereby making them a flexible information-gathering tool.

You can now open and edit active lists in grid views.

Active lists function differently than active channels. Active lists are not continuously re-evaluated and are not time-window constrained. Active lists draw from the event stream on the basis of their event or field/rule definitions and any rules designed to affect them.

You can use active lists as filters in other resources that are not based on active channels, such as reports.

In addition to their integral definitions, you can apply temporary (not saved) filters to active list grid views. Click the status description in the **Filter** line in the view header to use the Common Condition Editor.

Use the default items in the Active Lists resource tree for templates or for operational monitoring with minor modifications. For example, use the Trusted List to watch activity from known-to-be-safe IP sources and the Untrusted List to do the same for known unsafe sources.

If you have Administrator access you can have another group named All Active Lists that contains all active list groups and lists.



Note: For procedural information about working with active lists (including how to create, edit, delete, import, and export them), see ["List Authoring" on page 486](#).

Uses of Active Lists

The main uses of active lists are to:

- Maintain information, such as in the system content-provided Hostile List or Trusted List that maintain information on hostile and trusted IP addresses (and corresponding zones).
- Check for the existence of particular information in lists using the InActiveList condition (see ["Condition Tree Command Buttons" on page 890](#)).



Note: The InActiveList operator does not parse multi-value attributes. The InActiveList operator only evaluates single-value attributes, and treats multi-value attributes, such as Actor Account ID and Role, as single-value attributes.

For example, when a system is compromised (such as in a security breach), it can be added to the compromise list using rule actions. The information in the active list can then be used to collect all the events that occur on the asset while it is compromised. This can be used for tracking and further investigation on other systems that have come into contact with the compromised system

Active Lists for Long-Term State Retention

Active lists can store data over a longer period of time than rules or data monitors are capable of retaining. For example, rules can hold a state that describes the very recent past, normally few minutes. Data monitors may contain up to a day's worth of data, but data monitors usually contain aggregated data.

Active lists can answer the following question that cannot be addressed directly by rules or data monitors: "Has the source IP of the current event attacked one of my systems in the last 30 days?"

Optimize Data with Hash-Based Active Lists

A **hash-based active list** uses a hash function to map a set of data to a single number (a hash value).

To create a hash-based active list:

- Enable the **Optimize Data** option on the Active Lists **Attributes** tab.
- Make sure the list is set to **Case-Sensitive**

See ["List Authoring" on page 486](#).

The main advantage of using the hash-based active list by enabling the list's Optimize Data option is to reduce memory usage. Instead of storing the complete active list entry in memory, only the hash code (a number), count, and last modified time are stored. The complete entry is available in the database. Therefore, the size of each entry in memory is constant, regardless of the number of fields and corresponding data types in the active list schema.

In terms of performance, there is little or no difference between hash-based and regular active lists.

The Optimize Data option is useful for active lists that contains a large number of entries (for example, more than 100,000 entries) or a large amount of information per entry.



Note: There is a possibility of getting an inaccurate result from an active list that uses the Optimize Data option due to hash collisions. When two active list entries map to the same hash code, the result of the InActiveList condition can be inaccurate in some cases. However, the chances of two entries evaluating to same hash code are quite rare. In the current scheme, for an active list with 1 million entries, the chances of hash code contention are about 1 in 4,000,000.

You can switch active Lists between optimized (hashing) and non-optimized (non-hashing) after they are created.

Active List Monitor Events

The following monitor events include Active List Usage statistics. (See ["Status Monitor Events" on page 860](#) for more information.)

- Open Active Lists count (DEC: /Monitor/ActiveLists/ListCount)
- Active List entry count (DEC: /Monitor/ActiveLists/EntryCount)
- Active List entry capacity (DEC: /Monitor/ActiveLists/EntryCapacity)
- Active List entry usage (% of capacity used) (DEC: /Monitor/ActiveLists/EntryPercentUsed)
- Active List entry look-ups per second (DEC: /Monitor/ActiveLists/QueriesPerSecond)
- Active List entry updates per second (DEC: /Monitor/ActiveLists/ChangesPerSecond)
- Temporary Active Lists count (DEC: /Monitor/ActiveLists/TemporaryListCount)
- Temporary List entry count (DEC: /Monitor/ActiveLists/TemporaryEntryCount)

- Temporary Active List entry capacity (DEC: /Monitor/ActiveLists/TemporaryCapacity)
- Temporary Active List entry usage (% of capacity used) (DEC: /Monitor/ActiveLists/TemporaryPercentUsed)

Active Lists with Values

An active list with values divides the set of fields into key fields and value fields. Active lists with values provide the following functionality:

- Use an InActiveList condition to check the existence of an entry (using only keys or keys along with values). See ["Condition Tree Command Buttons" on page 890](#) for more about applying an InActiveList condition.



Note: The InActiveList operator does not parse multi-value attributes. The InActiveList operator only evaluates single-value attributes, and treats multi-value attributes, such as Actor Account ID and Role, as single-value attributes.

- Look up value fields for given key field values. Keys and values can consist of one or more columns.

A single key can map to a single value; for example, user name (key) to badge ID (value).

A single key can also map to multiple values (in a *multi-map* active list); for example, user name (key) to badge ID, first name, last name (three values). The actors resource is a typical use case for multi-map active lists. For example, you might map a single actor (key) to multiple roles (values).

Variables are used to retrieve the value portion of the active list entry.

To create an active list with values, select the **Fields-based** data option on the Active List editor Attributes tab, check **Key Fields** to enable a per-field Key option, and then select one or more data fields that must be unique. (For the complete procedure, see the topic on ["Creating or Editing an Active List" on page 486](#).)

Using Variables to Retrieve Data from Active Lists with Values

To add a list field in a condition, use the variable functions described in ["List Functions" on page 1075](#) that can yield list values.

When defining a variable to retrieve value information from active lists with values, be sure to specify these attributes for the variable:

- Name of the variable
- The active list to be used to retrieve values for the key
- Field mappings (mapping of the event fields to key fields in the active list)

For more about working with variables, see ["Variables" on page 1064](#).

Example: Active List with Values to Store Directory Information

As an example, suppose we want to create an active list with values to store directory information.

Create an Active List

We follow the basic procedure to create a new active list shown in ["Creating or Editing an Active List" on page 486](#). For the example, we create the active list with these options:

- Specify **Fields-based** data using **Key** fields
- The **Key** is the **Username**.
- The values contain various information corresponding to the given user name. For simplicity, you can store only user role information. (The user role usually determines the type of actions a user can take, and on what type of resources.) If desired, you can store additional information such as the user's First Name, Last Name, Phone Number, Email Address, and so on.

ActiveList : User Roles

Attributes Notes

+ Add Entry

Active List

* Name	User Roles
Optimize Data	☐
* Capacity (x1000)	10
* TTL Days	1
* TTL Hours	0
* TTL Minutes	0

(Name)
(Description)

* Data: ☐ Event-based ☒ Fields-based ☒ Key Fields

Name	Type	Sub-type	Key-field
Username	String		☒
User Roles	String		☐

OK Cancel Apply Help

Populate the Active List

We can populate the list in any of various ways:

- Manual data entry
- Export required information from Active Directory into a CSV file, and then import entries to the active list from the CSV file
- Use Active Directory User Group Puller tool
- Use event-based integration or other tools

Correlate Information Stored in UserRoles List

Once the Active Directory information is populated to an active list with values, we can access and correlate the user information using reports, rules, active channels, data monitors, and so on. The details of the correlation logic are as follows.

Create a Rule:

For this example, we choose a rule that does the following:

- Looks for events that update some critical database information
- Checks if the target user had privilege to perform the operation using the Active Directory User Role information, maintained in the active list

(For more information on creating rules, see ["Rules Authoring" on page 507](#) and ["Creating or Editing Rules" on page 509](#) .)

Use Variable to Get Role Information:

For the database update events, we can get the corresponding Active Directory role information using the **GetActiveListValue** variable.


The screenshot shows the 'Inspect/Edit' dialog box for the 'Global Variable:userRole'. The 'Parameters' tab is selected. The 'Function' dropdown is set to 'GetActiveListValue'. Below this, the text 'Retrieve ActiveList value' is displayed. The 'Arguments' section contains a 'List' dropdown set to 'User Roles'. The 'Field Mapping' section includes the instruction 'For each key field, select a matching SecurityEvent field.' and a table with three columns: 'Name', 'Field', and 'Key'. The table contains one row: 'Username' in the 'Name' column, 'Target User Name' in the 'Field' column, and a checked checkbox in the 'Key' column. The 'Preview' section has the instruction 'Set a value for each key field.' and a table with two columns: 'Name' and 'Value'. The 'Name' column has a checkbox labeled 'Key Fields' and the text 'Username' below it. The 'Value' column is empty. At the bottom right of the dialog is a 'Calculate' button. At the very bottom are 'OK', 'Cancel', 'Apply', and 'Help' buttons.

Name	Field	Key
Username	Target User Name	☒

Name	Value
☐ Key Fields Username	

Set Conditions to Check Role Permissions:

Once the role information is retrieved, we can check if the user has the role required to perform this operation. If the user does not have the required role, then the rule is triggered to alert the administrator to the unauthorized access.



Take Action Based on Results of Permissions Check:

If the user does not have required role, then rule can trigger and alarm the administrator regarding this unauthorized access. (This is configured on the Actions tab.)

Administrator

An administrator is a person who has the rights to administer ArcSight and manage users, groups, and their permissions.

See also:

- ["Users" on page 1063](#)
- ["User Types" on page 1063](#)
- ["Managing Users and Groups" on page 99](#)
- ["Managing Permissions" on page 107](#)


Advanced Editor

An Advanced Editor is available to accommodate special requirements based on context for providing input to various fields, conditions, or other values. The Advanced Editor provides different features, depending on the context in which it is called or used in the .


Typically, the Advanced Editor shows up during edit operations in the [Common Conditions Editor \(CCE\)](#), [Rules Editor](#) (for example, editing rule actions), and [Variables](#) editors, but it might show up in other contexts also.

Here are just a few examples of some of the contexts and features the Advanced Editor provides.

- Enable multi-line input. Advanced editors provide an option to enter multi-line input for rule action triggers, event field values in the common conditions editor (CCE), and so forth.



- Provide calendars and times.



Aggregation

Aggregation is a composition technique for building a new event from one or more existing events that support some or all of the new event's conditions.

You use aggregation to group occurrences of matching conditions based on incoming event field data values, and optionally count only distinct occurrences of those events. To support that, the provides Group By aggregation, in which you can group (aggregate) correlated events by field values. You can also optionally include distinct-value event processing combined with either join conditions and/or event grouping, to provide further constraints on when rules should fire.

Rules always run subject to their associated aggregation parameters, even if only the defaults. For more information, see ["Specifying Rule Thresholds and Aggregation" on page 528](#).

The *aggregated event count* is a derived event field available in the [Event Inspector](#), the [Common Conditions Editor \(CCE\)](#). The field also shows up in various data monitors, for example, [Moving Average Data Monitor](#). The aggregated event count is described with the [Event Group](#) data fields.

ArcSight Console

The is a graphical user interface that provides centralized intelligent real-time monitoring to secure your enterprise.

Console settings consist of your color selections, preferences, temporary filters, window sizes, and so on, and are saved in a .ast file. The current setting file you are using is displayed in the Console title bar (by default, machine:username.ast). You can perform operations to save or load .ast files stored locally, on the same machine where the Console is installed, on save or load .ast files stored and maintained by the Manager. After settings are saved in a file, the .ast file is listed in the File menu. The File menu lists the last four .ast files that have been accessed.

Concerning Console-Manager connections, you may want to note that while each Manager connects to many Consoles, each Console connects to only one Manager. Also, when a Console is connected to a Manager, it affects only that Manager, regardless of how that Manager may be linked within a larger Manager hierarchy.

If you are viewing various articles, information, reports, or command results in a Web browser, see the Support Matrix applicable to your ESM version for a list of supported browsers.

See related topics ["Working in the " on page 43](#).

Assets

Assets are network devices, installed throughout your enterprise, that you monitor for vulnerability or attack. Once asset information is stored, ESM tracks your assets and notifies you if they are exposed to a threat or vulnerability, or if they are attacked.

Within the Navigator panel's Assets resource tree there are a number of views of associated information. The Assets, Networks, Zones, Locations, Categories, and Vulnerabilities tabs each show different aspects of the devices in question.

When, how, and why you might need to modify the resources in the Assets tabs is described in Managing Assets and Associated Resources, and particularly in Changing Assets.

The shows assets, vulnerabilities, asset categories, scanner reports, and cases in active channels rather than static grid views. Leverage the power of channels for asset management, including use of filters, field sets, better sorting capabilities, and dynamic display of an unlimited number of items that is continually updated.

What the Assets tabs contain is described below.

Assets Tab

This view shows the population of your network as discreet entities with specific IP addresses and unique MAC and host names. You often use this view to pinpoint a particular asset, then double-click it to change its characteristics and associations in the Asset Editor. The presentation is hierarchical and shows only the assets to which you have access through the Asset Editor. Note that you can also identify mobile assets by MAC address.

Because the usage (Zones) and descriptive (Categories) views are separate, the Assets view is free to accurately describe the access restrictions that apply to a given user.

The Internet Address Range asset category has its own Asset Range Editor. The address range groups are standard spans of IP addresses provided as a convenience for your use in rules. You can also collectively reference these ranges using the named networks on the Zones tab. For example, a rule could reference the Dark Address Space item under System Zones to identify a category of source IP addresses from which traffic should not legitimately originate.

The Asset, Zone, and Asset Range Editors all include Categories and Zones tabs for editing these attributes.

The Asset Editor has an Alternate Interfaces tab. When a single device has multiple network interfaces, you can define each interface as an independent asset. Common examples of multiple interface devices are network connection points such as routers and bridges. To use

this editor from an appropriate asset in the tree, right-click it and choose **View Asset Alternate Interfaces**.

An asset or asset range (or its group) can belong to only one zone or location.

Zones Tab

The **Zones** tab shows the hierarchy of network-related logical (usage) groups into which assets are collected, and on which you can act through the Zone Editor.

You can also think of zones as aliases for portions of your network that are dedicated to certain organizational groups or functions. The Zones view can help disambiguate multiple private networks that might have overlapping address spaces.

When the zones in your enterprise are referencing multiple global or local zones, ArcSight networks can help disambiguate erroneous address space overlaps or gaps, especially for SmartConnectors. A zone or zone group can belong to only one network or location, and expresses a single contiguous address range.

Networks Tab

On this tab, you can view the hierarchical collection of network entities recognized within your system. In this context, a network is an enterprise-level registry of ArcSight zones. Networks are used to reconcile overlapping or missing asset ranges among zones (if they should erroneously occur). When networks are present, SmartConnectors use them to find their correct zone assignments. Note that networks apply only in enterprises that have networks broad enough to require multiple local or global maps. If your enterprise maps only its own address space (meaning that overlaps and gaps aren't likely) the Networks tab won't be populated.

Each Network resource can relate to only one Customer resource, but to multiple Zones, provided address ranges do not overlap.

Categories Tab

The hierarchy of asset categories provides a way to reference assets by means of their application or context. A given asset can be associated with multiple categories.

Asset categories are a cross-referencing capability that supports numerous business objectives. By making it possible to track network activity with certain assets on the basis of their business significance, data collection becomes possible; but just as importantly, other ArcSight analytical tools can also be brought to bear to derive many kinds of information. Finally, reporting capabilities can further analyze and permanently record the results.

The categories for a given enterprise are often quite specialized, but certain categories are usually present, even if customized.

Typical high-level asset categories often include:

- **ArcSight System Administration:** These assets include its Console, Databases, SmartConnectors, and Managers. The system automatically detects its own administrative assets and creates these entries in their respective groups.
- **Site Asset Categories:** These can be general monitoring categories such as Address Spaces, Applications, or Open Ports, but may include asset-tracking categories that are specific to issues such as business impacts and regulatory compliance. Business impacts might analyze the activity of a server that supports a particular product line. Regulatory compliance could monitor groups of workstations for HIPAA conformity.
- **System Asset Categories:** These can be any of a number of categorizations of assets such as Criticality (low, medium, high, very high), which would monitor by a classification of how crucial assets are to the enterprise.



Note: SmartConnector configuration also affects the ability to automatically create the assets that represent network devices. Each SmartConnector needs to report an IP address or hostname for its sensor so its events can be identified on the network. See the configuration guides for your SmartConnectors to ensure they are reporting this information.

Vulnerabilities Tab

The **Vulnerabilities** tab presents known vulnerabilities associated with the devices identified through the Assets tab.

Device vulnerabilities are presented as closely as possible with device descriptions to facilitate useful comparisons and easy reference

Locations Tab

The **Locations** tab shows the hierarchy of names your enterprise uses for its physical or geographical domains.

Similar to zones, you can think of locations as another type of alias. You use this alias for portions of your network that are referenced by where they are rather than by organizational group or function. A given asset or asset range can be associated with only one location, but a given location can be associated with any number of appropriate assets, zones, or their groups.

Attack

An exploited threat or an attempt to bypass security controls on a computer. The attack may alter, release, or deny data. Whether an attack succeeds depends on the vulnerability of the computer system and the effectiveness of existing countermeasures.

Audit Events

Audit events are events generated within the Manager to mark a wide variety of routine actions that can occur manually or automatically, such as adding an event to a case or when a Moving Average data monitor detects a rapidly rising moving average. Audit events have many applications, which can include notifications, task validation, compliance tracking, automated housekeeping, and system administration.

This topic lists the ArcSight audit events you can use in rules, filters, and other analytical or administrative resources. Observe the way these events are used in the standard system-related content for examples of how to apply them.

In the table [Audit Events on Resources](#), use the **Audit Event Category** to locate events. Use the Device Event Class (DEC) ID string in rules and filters. The **Audit Event Description** reflects the resource name you see in active channel grids. Additional details, when necessary, appear in the **Notes** column.

Compare audit events, which report on **system activity**, with Status Monitor events, which provide information about a wide variety of **system states**.

All resources (except actors, groups, and users) use the general audit events described in “Resources (Configuration Events Common to Most Resources),” in when a resource is added, deleted, updated, locked, or unlocked. Actors, groups, and users each use their own unique set of audit events. Other resources present unique audit events that are listed in this section in alphabetical order by resource.



Tip: To get *additional* details within the “update resource” audit events (beyond what is provided by default), you can enable a resource audit property called `resource.audit.update.uris` in the `file.server.properties` on the Manager to specify which resources should show extended audit event information.

For information on Logger audit events, see the Logger Administrator’s Guide appendix, “Logger Audit Events.”

Audit Events Common to Most Resources

These audit events are generated in response to creation events and configuration updates to most resources, except users, actors, and groups, which use different audit events. When a resource is added, deleted, updated, locked, or unlocked, the Manager generates one audit event with the following attributes:

- Device Event Class ID = resource:100 (deleted) or resource:101 (updated) or resource:102, and so on.
- Event Name = <resource type> deleted/updated/added.
- File Name = <Resource Name> (for example, John's Filter)
- File Path = <Resource URI> (for example, /All Filters/administrator's Filter/John's Filter)
- File Type = <Resource Type> (for example, Filter)

Audit Events on Resources

Audit Event Category	Device Event Class ID	Audit Event Description
Resource (Delete)	resource:100	Resource deleted. The Event Name describes the action and resource type (<ResourceName> deleted); for example, deleting a filter results in an event named Filter deleted.
Resource (Update)	resource:101	Resource updated. This audit event is generated when an existing resource is modified or added. See Resource (Add). The Event Name describes the action (update) and resource type (<ResourceName> updated); for example, modifying a report, results in an event name of Report updated.
Resource (Add)	resource:102	Resource added (inserted). The Event Name describes the action (insert) and resource type (<ResourceName> inserted); for example, adding a case, results in an event name of Case inserted. Adding a Case group results in an event name of Group [Case] inserted.
Resource (Lock)	resource:103	Resource locked <ResourceName> locked.
Resource (Unlock)	resource:104	Resource unlocked <ResourceName> unlocked.
	resourcereference:100	Could not locate a resource through the supplied universal resource identifier (URI).

Active Channel

Audit Events on Active Channels

Device Event Class ID	Audit Event Message
channel:001	An active channel [Channel Name] was opened/started.
channel:002	The channel [Channel Name] is empty; that is, there are no matching events for the built-in channel filter.
channel:003	The channel [Channel Name] query completed.
channel:004	The channel [Channel Name] query is slow.

Active List

Audit Events on Active Lists

Device Event Class ID	Audit Event Message
activelist:101	An entry was added to an active list.
activelist:102	An entry was removed from an active list.
activelist:103	An entry was changed in an active list.
activelist:104	An entry has expired in an active list.
activelist:105	An entry has been evicted from an active list. The active list is full and an entry is dropped.

Actor

Audit Events on Actors

Device Event Class ID	Audit Event Message
actor:100	An actor was deleted.
actor:102	An actor was created.
actor:110	One or more actor attributes were updated.
actor:111	A multi-valued actor attribute was added to an actor.
actor:112	A multi-valued actor attribute was removed from an actor.

Archive

Audit Events on Archive Management

Device Event Class ID	Audit Event Description
archive:100	Archive created
archive:101	Archive deleted
archive:102	Event archive settings updated
archive:110	Archive activated
archive:111	Archive activation cancelled
archive:112	Archive activation failed
archive:120	Archive operation succeeded
archive:121	Archive operation cancelled
archive:122	Archive operation failed
archive:130	Archive deactivated
archive:131	Archive deactivation cancelled
archive:132	Archive deactivation failed
archive:140	Archive scheduled
archive:141	Archive schedule cancelled
archive:142	Archive schedule failed
archive:143	Supplemental archive operation succeeded
archive:144	Supplemental Archive — No data
archive:145	Supplemental archive, failure

Authentication

Audit Events on Authentication

Audit Event Category	Device Event Class ID	Audit Event Message
Authentication	authentication:100	A client authenticated with the Manager.
Authentication	authentication:101	A client authentication login failed.
Authentication	authentication:102	An authenticated client logged out of the Manager.
Authentication	authentication:103	Authentication logout time.

Audit Events on Authentication, continued

Audit Event Category	Device Event Class ID	Audit Event Message
Authentication	authentication:104	A client made several unsuccessful attempts to log in to the Manager, resulting in an excessive number of failed logins.
Authentication	authentication:105	A non-FIPS client authenticated with the Manager via login. (A valid login by a non-FIPS ArcSight Console authenticating itself to the Manager triggers this audit event.) For information on how to configure a non-FIPS client (such as ArcSight Console) to log in to a FIPS-enabled Manager, see the ESM Administrator's Guide.
Connector Login	authentication:200	Successful connector authentication.
Connector Login	authentication:201	Connector authentication failed.
Authentication	authentication:202	A non-FIPS connector authenticated with the Manager via login. (A valid login by a non-FIPS SmartConnector authenticating itself to the Manager triggers this audit event.) For information on how to configure a non-FIPS SmartConnector to connect to a FIPS-enabled Manager, see the document, <i>Installing FIPS-Compliant SmartConnectors</i>.

Authorization

Device Event Class ID	Audit Event Description
authorization:100	Manager refused to authorize client.

Backpressure Audit Events

Backpressure audit events can be viewed by creating an active channel with the filter Device Event Class ID StartWith "backpressure".

Backpressure Audit Events

Device Event Class ID	Audit Event Message
backpressure:100	Automatic backpressure is on
backpressure:101	Automatic backpressure is off
backpressure:102	Manual backpressure is on
backpressure:103	Manual backpressure is off
backpressure:104	Acceptable lag was set to <acceptable_lag>

Backpressure Audit Events, continued

Device Event Class ID	Audit Event Message
backpressure:105	Backpressure mode was set to <mode>
backpressure:106	Partition <id> is triggering backpressure
backpressure:107	Partition <id> is no longer triggering backpressure

Connector Connection



Note: Audit events with DEC IDs beginning with `agent` indicate statistics collected at the connector level. These audit events are different from those described in [SmartConnector Flow Statistics](#), which are event flow from connectors monitored at the Manager level.

Audit Events on Connector Connections

Device Event Class ID	Audit Event Description
agent:009	Manager rejected a connection attempt from a connector for reasons other than authentication failure.
agent:030	Connector started.
agent:031	Connector shutting down.
agent:041	Event flow on connector has been started, stopped, or paused.
agent:043	This occurs at the specified interval, if Device Status Monitoring is enabled. It Reports the last time an event was received and the number of events since the connector started and since the previous report.
agent:101	Connector has just connected to Manager.
agent:102	Connector is sending events but no heartbeats.
agent:103	Connector is sending neither events nor heartbeats.
agent:104	An unknown connector attempted to connect to the Manager.
agent:105	A connector presented an incorrect shared secret when authenticating.

Connector Exceptions



Note: Audit events with DEC IDs beginning with `agent` indicate statistics collected at the connector level. These audit events are different from those described in [SmartConnector Flow Statistics](#), which are event flow from connectors monitored at the Manager level.

Audit Events on Connector Exceptions

Device Event Class ID	Audit Event Description
agent:012	Connector detected source events from a sensor device containing incorrect time stamps.
agent:013	Connector noted that a new sensor device is sending events.
agent:014	Connector could not find a base event referenced in a syslog aggregate event.
agent:015	connector connection device failure.
agent:016	connector connection device success.
agent:017	Connector successfully executed a command.
agent:018	Connector could not execute a command.
agent:019	Connector is caching events because they could not be immediately transmitted to the Manager.
agent:020	Connector has emptied its cache of events.
agent:021	Connector could not communicate with an NT collector sensor.
agent:023	Connector could not communicate with a CheckPoint sensor.
agent:024	Connector is having difficulty communicating with CheckPoint.
agent:028	Connector experienced an unexpected problem.
agent:029	Connector was forced to drop its cached data.
agent:035	Connector sent an event with a bad timestamp; it is beyond the retention period.
agent:053	Connector starts dropping events/files from agent queue. - deviceEventCategory: /Agent/Queue/Dropped - agentSeverity: High (3)
agent:054	Connector has stopped dropping a number of events from queue files. - Name: Agent queue dropped events - deviceEventCategory : /Agent/Queue/Events/Dropped - agentSeverity: High (3)

Audit Events on Connector Exceptions, continued

Device Event Class ID	Audit Event Description
agent:055	Connector has stopped dropping a number of files from Agent Queue. - Name: Agent queue dropped files. - deviceEventCategory: /Agent/Queue/Files/Dropped - agentSeverity: High (3)
agent:107	Connector is generating events for too many devices. This may be due to improper field mapping in the connector that causes what should be the same device to be referenced in too many different ways. The correlated internal event sets these additional fields: - Name: Too many devices being created - possible parsing problem - deviceEventCategory: /Agent/SideTableOverflow - agentSeverity: Very high (4)

Connector Login



Note: Audit events with DEC IDs beginning with `agent` indicate statistics collected at the connector level. These audit events are different from those described in [SmartConnector Flow Statistics](#), which are event flow from connectors monitored at the Manager level.

Audit Events on Connector Logins

Device Event Class ID	Audit Event Description
authentication:200	Successful connector authentication.
authentication:201	Connector authentication failed.

Connector Registration and Configuration



Note: Audit events with DEC IDs beginning with `agent` indicate statistics collected at the connector level. These audit events are different from those described in [SmartConnector Flow Statistics](#), which are event flow from connectors monitored at the Manager level.

Audit Events on Connector Registration and Configuration

Device Event Class ID	Audit Event Description
agent:007	Connector successfully registered with Manager.
agent:008	Connector did not successfully register with Manager.

Audit Events on Connector Registration and Configuration, continued

Device Event Class ID	Audit Event Description
agent:010	Connector upgrade succeeded. This is currently in the context of an installer upgrade.
agent:011	Connector upgrade failed. This event is not currently being generated.
agent:022	Connector could not process a reconfiguration request.
agent:025	Connector content was successfully updated.
agent:026	Connector content update failed.
agent:032	Connector configuration was successfully changed.
agent:050	Used to track EPS consumption statistics

Content Management

Audit Events on Content Management

Device Event Class ID	Audit Event Description
contentmanagement:100	ESM Manager is now a content publisher.
contentmanagement:102	ESM Manager is now a subscriber.
contentmanagement:103	ESM Manager is no longer a subscriber.
contentmanagement:120	A package has been manually delivered to subscribers.
contentmanagement:121	A package has been received from the content publishing system.
contentmanagement:122	A package delivery has started.
contentmanagement:123	A package delivery has completed.
contentmanagement:124	A package delivery attempt has failed.
contentmanagement:125	A package has been successfully installed by a subscriber.
contentmanagement:126	A package has failed to install for a subscriber.
contentmanagement:130	A package has been added to a schedule.
contentmanagement:131	A package has been removed from a schedule.
contentmanagement:140	A content synchronization schedule has been enabled.
contentmanagement:141	A content synchronization schedule has been updated.
contentmanagement:142	A content synchronization schedule has been disabled.
contentmanagement:143	An attempt to enable or update the content synchronization schedule has failed.

Created Ticket on External System

Audit Events on External System

Audit Event Message	Device Event Category	Device Event Class ID
Created Ticket on External System	/externalticket/snow	externalticket:100

Dashboard

Audit Events on Dashboards

Device Event Class ID	Audit Event Description
dashboard:001	A data monitor on a dashboard was newly accessed after not having been accessed for some time (for example, the dashboard had been closed). This audit event is generated on a per-user, per-Console-session basis.
dashboard:100	Dashboard has opened.

Data Monitors

Audit events related to data monitors are described in the following tables, categorized by data monitor type. (See also the Dashboard audit events topic.)

All Data Monitors

Audit Events on all Data Monitors

Device Event Class ID	Audit Event Description
<code>datamonitor:001</code>	This audit event contains the following values: • The count of events processed by a data monitor. • The elapsed time (in seconds) that a data monitor spent processing those events. The values given above are measured starting from the time that the last event of this type was generated for the given data monitor.
<code>datamonitor:002</code>	This audit event contains the following values: • The count of distributed cache synchronization points for a data monitor. • The elapsed time (in seconds) that a data monitor spent processing those synchronization events. The values given above are measured starting from the time that the last event of this type was generated for the given data monitor. This event is only generated if ESM is running in distributed mode. In compact mode, there is no distributed cache, and therefore there are no synchronization points.

The audit events mentioned above are also called **Data Monitor Telemetry** events. Each event is a measurement of the values given above since the last event of this type was generated for the given data monitor. The internal counters are set to zero after each event is generated. The result is that each event will contain the count and elapsed processing time for the given data monitor over the scheduled telemetry interval, which is 5 minutes by default. For example, if the telemetry event reports that a given data monitor processed 1000 events and the elapsed time was 0.25 seconds, this means that the data monitor spent a total of 0.25 seconds to process the 1000 events over a 5 minute interval.

The audit events mentioned above are designed to be useful in **Top Value Counts** data monitors. They are used by the **Data Monitor Status** dashboard. For more information on the **Data Monitor Status** dashboard, see the [ArcSight Administration and ArcSight System Standard Content Guide](#).

Last State Data Monitors

Audit Events on Last State Data Monitors

Device Event Class ID	Audit Event Description
datamonitor:400	A Last State data monitor entry has exceeded its time-out period and was automatically removed.
datamonitor:401	A Last State data monitor entry value was manually changed by the user.
datamonitor:402	A Last State data monitor entry was manually removed by the user.

Moving Average Data Monitor

Audit Events on Moving Average Data Monitors

Device Event Class ID	Audit Event Description
datamonitor:101	Moving average threshold.
datamonitor:102	Moving Average data monitor detected a rapidly falling moving average
datamonitor:103	Moving Average data monitor detected a rapidly rising moving average.
datamonitor:104	Moving Average data monitor reporting the current moving average.
datamonitor:105	A value was added to a Moving Average data monitor, which is now monitoring a new Group-By set of values.
datamonitor:106	A value was removed from a Moving Average data monitor. The data monitor is no longer monitoring a particular Group-By set of values.

Reconciliation Data Monitor

Device Event Class ID	Audit Event Description
datamonitor:300	Correlation data monitor reporting a correlated or non-correlated event.

Statistical Data Monitor

Audit Events on Statistical Data Monitors

Device Event Class ID	Audit Event Description
datamonitor:200	Statistical Data Monitor reported a change in status.
datamonitor:201	A value was added to a Statistical Data Monitor, which is now monitoring a new Group-By set of values.
datamonitor:202	A value was removed from a Statistical Data Monitor. The data monitor is no longer monitoring a particular Group-By set of values.

Top Value Counts Data Monitor

Audit Events on Top Value Counts

Audit Event Category: Moving Average Data Monitor	
Device Event Class ID	Audit Event Description
datamonitor:500	For a Top Value Counts Data Monitor, the top N counts (N events).
datamonitor:501	Counts that were most recently added to the data monitor (from 0 ... N events).
datamonitor:502	Counts that were most recently removed from the data monitor (from 0 ... N events).

Distributed Correlation

The audit events described in this topic pertain to ESM that is installed in a distributed correlation mode, not compact mode. Audit events are categorized according to the services within a cluster: aggregator, correlator, dcache, mbus, and persistor.

For related information, refer to the following ESM publications:

- [ESM Installation Guide](#) for instructions on installing ESM in distributed mode
- [ESM 101](#) for distributed correlation concepts
- [ESM Administrator's Guide](#) for details on configuring and managing services in your distributed correlation cluster

Aggregator Audit Events

Audit Events on the Aggregator Service

Device Event Class ID	Audit Event Description
aggregator:100	Aggregator is up.
aggregator:101	Aggregator is down.
aggregator:103	Messages received in aggregator.

For rule aggregation concepts, see the topic, "Rule Aggregation," in *ESM 101*.

To define rule thresholds and aggregation over a time criteria, see ["Specifying Rule Thresholds and Aggregation" on page 528](#).

Correlator Audit Events

Audit Events on the Correlator Service

Device Event Class ID	Audit Event Description
correlator:100	Correlator is up.
correlator:101	Correlator is down.
correlator:103	Event count received in correlator.

For correlation and event matching concepts, see the topic "Correlation Overview" in *ESM 101*.

To define rule conditions, see ["Specifying Rule Conditions" on page 513](#).

DCache (Distributed Cache) Audit Events

The DCache service manages short-term storage of data needed for internal cluster operations.

Audit Events on Distributed Cache (DCache) Service

Device Event Class ID	Audit Event Description
dcache:100	DCache connection is up.
dcache:101	DCache connection is down.

MBus (Message Bus) Audit Events

The message bus control and message bus data service handles messaging among cluster components.

Audit Events on MBus Service

Device Event Class ID	Audit Event Description
mbus:100	MBus connection is up.
mbus:101	MBus connection is down.
mbus:102	Message count remaining in MBus.
mbuscontrol:100	Message bus control is up.
mbuscontrol:101	Message bus control is down.
mbusdata:100	Message bus data is up.
mbusdata:101	Message bus data is down.
mbus:104	Message bus events consumption pause.
mbus:103	Message bus events consumption resumed.

Persistor Audit Events

Persistor statistics refer to the persistor service in a distributed correlation environment. The persistor service writes to disk the information that needs to be retained, retrieved, and shared.

Persistor statistics are maintained for licensing purposes.

Status Monitor Event Categories for Persistor Statistics

Device Event Class ID	Description
monitor:104	Total EPS
monitor:109	EPS rate at GB/day

Distributed Event Forwarding

Distributed event forwarding produces audit events that provide visibility into how many events have been forwarded in a given time period. This audit event is produced once per minute per correlator. For example, if there are three correlators then three of these events are produced per minute. The total number of events forwarded in that time period can be derived by summing the event counts from each event.

Device Event Class ID	Audit Event Description
forwarder:001	Count of events forwarded to destination The event field contains the following values: - Count of events forwarded by this correlator in this time period - Elapsed time (in seconds) of this time period - ID of the correlator that produced this event (for example, "correlator1")

Global Variables

The following events also apply to resources in general. See ["Audit Events Common to Most Resources" on page 834](#).

Audit Events on Global Variables

Device Event Class ID	Audit Event Description
resource:100	Global variable deleted.
resource:101	Global variable updated.

Audit Events on Global Variables , continued

Device Event Class ID	Audit Event Description
resource:102	Global variable inserted.
resource:103	Global variable locked.
resource:104	Global variable unlocked.

Group Management

The following audit events are generated for any group add, update, or delete, including user groups. The details of the which type of resource was configured or modified are provided in the event name. (For more information on user management audit events, see the User Management category.)

Audit Events on Group Management

Device Event Class ID	Audit Event Description
group:100	A group was deleted.
group:101	A group was updated. This audit event is generated when an existing group is modified or added.
group:102	A group was added (group inserted). When a new group is added, two audit events are generated: this event (group:102), and a Group Update audit event (group:101).

License Audit

Each of these events is reported every 24 hours, beginning 24 hours after you start the Manager.

Audit Events on License Audit

Device Event Class ID	Audit Event Description
license:100	The number of assets you have at this time.
license:101	The number of devices you have at this time.
license:102	The number of actors you have at this time.
license:103	The number of Console users you have at this time.
license:104	The number of web users you have at this time.

Audit Events on License Audit , continued

Device Event Class ID	Audit Event Description
license:105	The average number of incoming events per second (EPS) over the last 24 hours and whether it exceeds your license.
license:106	The number of times that event-105 threshold breaches have occurred since the Manager started, and the license limit.
license:107	The number of times that EPS violations have breached the threshold over the number of days specified in your license. This is a serious license violation. For more information look at the License:105 and License:106 events.

Logger Component

ArcSight Logger is a component in CORR-Engine. The available audit events for this component vary by product version; some events in this topic may not apply to your version.

Logger component audit events are stored in the Default Storage Group. As a result, these events can be searched using the ArcSight Command Center search. For example, you can search for this event: `"/Logger/Resource/Archive/Configuration/Add/".`

A Logger audit event contains information about the following prefix fields:

- Device Event Class ID (DEC ID)
- Device Severity
- Message
- Device Event Category—(keyName for this CEF extension is “cat”)

For example:

```
Sep 19 08:26:10 zurich CEF:0|ArcSight|Logger|3.5.0.13412.0|logger:500|Filter added|2| cat=/Logger/Resource/Filter/Configuration/Add msg=Filter [Unified Query Test] has been added
```

The following tables list the information contained in audit events generated by the Logger component in ESM with CORR-Engine. Where applicable, the tables also show additional fields. The Severity for all Logger component events is 2.

Alerts

Audit Events on Alerts

Device Event Class ID	Audit Event Description	Additional Fields
logger:610	Alert [name] has been added	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmplpAddr dvchost=syslogOrSnmplpHost HostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmplpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger: 611	Alert [name] has been deleted	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmplpAddr dvchost=syslogOrSnmplpHost HostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmplpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses

Audit Events on Alerts, continued

Device Event Class ID	Audit Event Description	Additional Fields
logger:612	Alert [name] has been updated	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmplpAddr dvchost=syslogOrSnmpHostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger:613	Alert [name] has been enabled	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmplpAddr dvchost=syslogOrSnmp HostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses

Audit Events on Alerts, continued

Device Event Class ID	Audit Event Description	Additional Fields
logger:614	Alert [name] has been disabled	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmplpAddr dvchost=syslogOrSnmp HostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger:615	Alert [name] has been sent	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmplpAddr dvchost=syslogOrSnmpOr EsmHostName cn1Label=SyslogoOrSnmpOr EsmDestination Port cn1=syslogOrSnmpOrEsmPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses

Archives

Audit Events on Archives

Device Event Class ID	Audit Event Description	Additional Fields
logger:520	Archive [archiveName] has been added	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:521	Archive [archiveName] has been deleted	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:523	Archive [archiveName] has been loaded	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:524	Archive [archiveName] has been unloaded	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:525	Archive [archiveName] has been archived	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId

Audit Events on Archives, continued

Device Event Class ID	Audit Event Description	Additional Fields
logger:526	Event archive settings added	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:527	Daily archive task settings updated	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger:528	Event archive failed	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId

Certificates

Audit Events on Certificates

Device Event Class ID	Audit Event Description	Additional Fields
logger:643	Certificate [name] has been added	fname=alias duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate
logger:650	Certificate [name] has been deleted	fname=alias duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate

Audit Events on Certificates, continued

Device Event Class ID	Audit Event Description	Additional Fields
logger:651	Certificate [name] has been updated	fname=alias duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate

Peers

The peers covered in this topic refer to peer relationships between or among the Logger components in multiple installations of ESM with CORR-Engine.

Audit Events on Logger Peers

Device Event Class ID	Audit Event Description	Additional Fields
logger:550	Peer Logger [name] has been added	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger fileId=LoggerId
logger:570	Peer Logger authorization [name] has been added	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger Authorization
logger:571	Peer Logger authorization [name] has been deleted	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger Authorization fileId=LoggerId

Saved Searches

Audit Events on Saved Searches

Device Event Class ID	Audit Event Description	Additional Fields
logger:540	Saved search [name] has been added	fname=savedSearchName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
logger:541	Saved search [name] has been deleted	fname=savedSearchName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
logger:542	Saved search [name] has been updated	fname=savedSearchName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId

Searches

Audit Events on Searches

Device Event Class ID	Audit Event Description	Additional Fields
logger:680	Search indices have been added or Search index has been added	cs4=sessionId fileType=Search Index Configuration duser=UserName msg=Search index has been added cn1=1 duid=1 cs4Label=Session ID rt=receiptTime cn1Label=No. of fields added
logger:690	Search options have been updated	cs6=false cs7=true cs4=sessionId cs5=false cs2=false cs3=false cs1=true cs8=false cs1Label=Field Search Case Sensitivity duid=1 cs7Label=Field Summary cs8Label=Field Summary Field Discovery cs6Label=Display options rawEvent cs3Label=Regex Search Unicode Case Sensitivity fileType=Search Options duser=UserName cs5Label=Regex Search Canonical Equality Check cs4Label=Session ID rt=receiptTime cs2Label=Regex Search Case Sensitivity
logger:710	Search session [sessionId] has been cancelled by [user]	cs1Label=Session ID duid=1 cs1=sessionIdfile duser=UserName rt=receiptTime

Search Filters

Audit Events on Filters

Device Event Class ID	Audit Event Description	Additional Fields
logger:500	Filter [filterName] has been added	fname=filterName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId
logger:501	Filter [filterName] has been deleted	fname=filterName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId
logger:502	Filter [filterName] has been updated	fname=filterName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId

Storage Groups

Audit Events on Storage Groups

Device Event Class ID	Audit Event Description
logger:530	Storage group [storageGroupName] has been added
logger:532	Storage group [storageGroupName] has been updated

Storage Volume

Audit Event on Storage Volume

Device Event Class ID	Audit Event Description	Additional Fields
logger:536	Storage volume [name] has been added	fname=storageVolumeName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Volume fileId=storageVolumeld

Manager Activation

Device Event Class ID	Audit Event Description
manager:100	Manager has started.
manager:101	A clean Manager shutdown has been requested.

Manager External Event Flow Interruption

Device Event Class ID	Audit Event Description
manager:200	Manager has stopped the event flow.
manager:201	Manager has allowed the event flow to resume.

Mark Similar

The creation or deletion of mark similar configurations generates audit events described below. You can add filters to view these events.

Device Event Class ID	Message	Priority
marksimilar:102	Mark similar configuration created	Low
marksimilar:100	Mark similar - all have been removed	Medium
marksimilar:100	Mark similar configuration removed due to time window expiry	Low
marksimilar:100	Mark similar configuration removed due to error. Check server.log	High

Status Monitor Events

The status monitor events are internal ESM events that appear on the event stream. These events can reveal and isolate many different quantity and time-unit issues that bear directly on performance and capacity. There are many possible applications of this system-state data, but those applications must always be interpreted within the context of your particular hardware, software, and network environment, and the deployment choices you have made.

Compare status monitoring events, which provide information about a wide variety of **system states**, to Audit Events, which report on **system activity**.

ESM does not provide standard content (such as filters, active channels, reports, query viewers, and so on) on the status monitors, but if you are interested in all /Monitor events, create an active channel and use the condition `Device Event Category StartsWith /Monitor`, or use a very specific condition, for example, `Device Event Category = /Monitor/Asset/TotalCount`. You can also use the condition in a query for a report or a query viewer.

Active Channel Statistics

Active channel statistics, specifically any changes that occur in the counts they report, can indicate performance issues and the use of processing cycles. These events summarize:

- The number of currently open Active Channels
- The number of events inserted into Active Channels per second
- The number of events changed across all open Active Channels per second

Status Monitor Event Categories for Active Channels

Device Event Class ID	Audit Event Description	Notes
monitor:100	Open active channel count	Count, current value
monitor:174	Active channel event insertions per second	Count per second, since last monitor event
monitor:175	Active channel event changes per second	Count per second, since last monitor event

Active List Statistics

Active list statistics monitor the resources being used by active lists. Active lists entries use some memory and database resources, and use CPU resources when they are referenced by other parts of the system (for example, rules, reports, and filters). While changes to these

temporary lists are not persisted, they do represent some memory overhead. Note that when active lists are used by replay-with-rules, this also creates temporary lists.

Status Monitor Event Categories for Active List Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:114	Open active list count	Count, current value
monitor:115	Active list entry count	Count, current value
monitor:116	Active list entry capacity	Count, current value
monitor:117	Active list entry usage	Percent, current value
monitor:118	Temporary Active list count	Count, current value
monitor:119	Temporary Active list entry count	Count, current value
monitor:120	Temporary Active list capacity	Count, current value
monitor:121	Temporary Active list usage	Percent, current value
monitor:122	Active list queries per second	Count per second, since startup
monitor:123	Active list changes per second	Count per second, since startup

Asset Statistics

Asset statistics offer insight into performance areas that affect assets in the system and can help resolve source, destination, agent, and device asset issues for incoming events. These events summarize:

- **Asset resolutions per second** is the average number of end-points in events, that are resolved to assets in a second.
- **Asset resolutions average time** is the average time in milliseconds taken to resolve an end-point in an event to an asset.
- **Asset scanner events per second** is the number of scanner events processed in a second.
- **Asset scanner events average time** is the average time in milliseconds taken to process a scanner event.

Status Monitor Event Categories for Asset Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:200	Asset total count	Count, current value
monitor:201	Scanner events processed per second	Count per second, since last monitor event
monitor:202	Asset resolutions per second	Count per second, since last monitor event
monitor:203	Scanner event average processing time	Count per second, since startup

Status Monitor Event Categories for Asset Statistics, continued

Device Event Class ID	Audit Event Description	Notes
monitor:204	Asset resolution average time	Microseconds per count, since startup
monitor:205	Asset source resolution average time	Microseconds per count, since startup
monitor:206	Asset destination resolution average time	Microseconds per count, since startup
monitor:240	Transitive closure size	Count, current value

Data Monitor Statistics

The data monitor statistics indicate how intensively the data monitors are working, which in turn can indicate situations such as filters needing adjustment or data monitors needing restructuring. These events summarize:

- **Active probes** is the number of currently enabled data monitors.
- **Evaluations per second** is the number of events times the number of enabled data monitors per second.

Status Monitor Event Categories for Data Monitor Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:101	Active data monitor probe count	Count, current value
monitor:124	Data monitor evaluations per second	Count per second, since last monitor event

Transformation Hub Statistics

These statistics monitor is for reading events from, and writing events to, the database. As such, they are database health indicators. These events summarize:

- **Event count** is the number of events inserted into the database since the last monitor event.
- **Insert time** is the average time taken to insert each event into the database, in microseconds.
- **Retrieval time** is the average time taken to retrieve each event from the database in microseconds.

Status Monitor Event Categories for Transformation Hub Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:102	Events insertion time per event	Microseconds per count, since last monitor event
monitor:103	Events processed count	Count, since last monitor event
monitor:140	Events retrieval time per event	Microseconds per count, since last monitor event

Filter Engine Statistics

The count of in-memory filter evaluations can serve as a broad indicator of filter performance.

Status Monitor Event Category for Filter Engine Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:161	Filter evaluation count	Count, since last monitor event

Main Flow Statistics

These events report statistically on the overall throughput of the Manager, for both incoming and internal events. This flow is the sequence of processing steps applied to each event and is a broad indicator or benchmark of system traffic. These events summarize:

- **Count** describes the number of events that have passed through the flow since the manager started.
- **Rate** describes the current event rate in events per second.

Status Monitor Event Categories for Main Flow Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:230	Main flow event rate	Count per second, since last monitor event
monitor:231	Main flow event count	Count, since startup

Notification Statistics

This group reports on notification activity, which can be of diagnostic value in detecting unusually high notifications activity.

- **New count** describes the number of new notifications since the last monitor event.
- **Escalated count** describes the number of notifications that were escalated since the last monitor event.

Status Monitor Event Categories for Notification Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:180	New notification count	Count since last monitor event.
monitor:181	Escalated notification count	Count since last monitor event.

Threat Detector Statistics

These events provide statistics for recent or pending threat detection runs. Because threat detection is database-intensive, these statistics can indicate or help diagnose database performance issues.

Status Monitor Event Categories for Threat Detector Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:190	Pattern discoveries run count	Count, since last monitor event.
monitor:191	Pattern discoveries queued count	Count, current value.

Report Statistics

These events provide statistics about the current number of reports querying the database or being rendered. Because reports are database-intensive, these statistics can indicate or help diagnose database performance issues.

Status Monitor Event Categories for Report Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:130	Reports running count	Count, current value.
monitor:131	Reports querying database count	Count, current value.
monitor:132	Reports rendering count	Count, current value.

Resource Framework Statistics

Resource-framework events report on the database activity connected with updates (reads, writes, and deletions) to system resources such as rules, assets, and filters, since the last monitor event. This data can be valuable in tracking or diagnosing performance-related issues such as automatic asset maintenance, the threat-level formula, or rule-driven usage.

Status Monitor Event Categories for Resource Framework Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:171	Resources inserted per second	Count per second, since last monitor event.
monitor:172	Resources updated per second	Count per second, since last monitor event.
monitor:173	Resources deleted per second	Count per second, since last monitor event.

Rules Engine Statistics

The statistics related to the Manager's rules engine can help reveal performance issues in several areas. Remember that information about rules activity always needs to be considered in the full content of the Manager's operations. For example, a busy Moving Average data monitor, if used inefficiently, can affect several of these statistics; a poorly written rule can inadvertently drive up the rate of actions executed.

These statistics have the following performance implications

- Count of events inserted into the rule engine: CPU.
- Rate of event insertion into the rule engine: CPU.
- Count of correlated events generated by the rule engine: CPU.
- Rate of correlated event generation by the rule engine: CPU.
- Count of events that are still present in rule engine's working memory: memory.
- Count of groupBy cells that are being used by the rule engine: memory.
- Count of rules currently active in the rule engine: comparative value only.
- Rate of actions being executed by the rule engine: CPU.
- Count of events matching any rule: CPU, memory.
- Count of events matching a rule with single alias: CPU, memory.
- Count of events matching a rule with multiple aliases: CPU, memory.
- Count of events rule matches: CPU, memory.

Status Monitor Event Categories for Rules Engine Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:151	Rules total event count	Count, since last monitor event.
monitor:152	Rules inserted events per second	Count per second, since last monitor event.
monitor:153	Rules generated events per second	Count per second, since last monitor event.
monitor:155	Rules in-memory event count	Count, current value.
monitor:156	Rules group by cells size	Count, current value.
monitor:157	Active rules count	Count, current value.
monitor:158	Rules actions rate	Count per second, since last monitor event.
monitor:159	Rules generated event count	Count, since last monitor event.
monitor:232	Events matching any rule	Count, since last monitor event.

Status Monitor Event Categories for Rules Engine Statistics, continued

Device Event Class ID	Audit Event Description	Notes
monitor:233	Events matching filter rule	Count, since last monitor event.
monitor:234	Events matching join rule	Count, since last monitor event.
monitor:235	Match Count	Count, since last monitor event.

Session List Statistics

Session list statistics monitor the resources being used by session lists. Session lists entries use some memory and database resources, and use CPU resources when they are referenced by other parts of the system (for example, rules, reports, and filters).

Status Monitor Event Categories for Session List Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:260	Open session list count	Count, current value.
monitor:261	Session list entry count	Count, current value.
monitor:262	Session list entry capacity	Count, current value.
monitor:263	Session list entry usage	Percent, current value.
monitor:264	Session list queries per second	Count per second, since startup.
monitor:265	Session list changes per second	Count per second, since startup.

Session Management Statistics

This statistic tracks the current number of active user sessions.

Status Monitor Event Category for Session Management Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:160	Active session count	Count, current value.

SmartConnector Flow Statistics

SmartAgent flow statistics record the event rates that occur at different stages of agent processing. "Sum of" statistics are sums of all values reported by all agents connected to the Manager. All values are statistics over the past 1-minute range. These events summarize:

- **Received event rate** is the rate at which agents receive events from devices.
- **Post filter event rate** is the rate of events that passed the filter (for example, were not filtered out).

- **Post aggregation event rate** is the rate of event aggregation.
- **Agent-to-manager event rate and count** describe how many events were actually sent to the Manager.
- **Cache size** describes the estimated size of the on-disk agent event cache.

Status Monitor Event Categories for SmartConnector Flow Statistics

Device Event Class ID	Audit Event Description	Notes
monitor:104	Agent to-manager event count	Count, since startup.
monitor:109	Agent to-manager event rate	Count per second, since last monitor event.
monitor:110	Agent received event rate	Count per second, since last monitor event.
monitor:111	Agent post-filter event rate	Count per second, since last monitor event.
monitor:112	Agent post-aggregation event rate	Count per second, since last monitor event.
monitor:113	Estimated agent cache size	Count, current value.
monitor:141	Sum of agent to-manager event counts	Count, since startup.
monitor:146	Sum of agent to-manager event rates	Count per second, since last monitor event.
monitor:147	Sum of agent received event rates	Count per second, since last monitor event.
monitor:148	Sum of agent post-filter event rates	Count per second, since last monitor event.
monitor:149	Sum of agent post-aggregation event rates	Count per second, since last monitor event.
monitor:150	Sum of estimated agent cache sizes	Count, current value.

Notification

Audit Events for Notification Category

Device Event Class ID	Audit Event Description
notification:100	Notification has been disabled.
notification:101	Notification has been disabled because the queue of notifications to be sent is too large.
notification:102	Notification has been enabled.

Audit Events for Notification Category, continued

Device Event Class ID	Audit Event Description
notification:103	Notification has been enabled because the queue of notifications is back under control.
notification:104	A particular notification destination has been disabled.
notification:105	A particular notification destination has been disabled because too much traffic was directed at it.
notification:106	A particular notification destination has been enabled.
notification:107	A notification expired without being acknowledged.
notification:108	A functioning destination could not be located for this notification.
notification:109	An old notification has been purged.

Notification Acknowledgement, Escalation, and Resolution

Device Event Class ID	Audit Event Description
notification:110	Notification has been escalated.
notification:111	Notification sent requires acknowledgement.
notification:112	An informational notification was sent.
notification:300	This notification has been acknowledged.
notification:301	This notification has been resolved.

Notification Testing

Device Event Class ID	Audit Event Description
notification:200	Sent a test notification to this destination group.

Threat Detector

Device Event Class ID	Audit Event Description
pattern:001	New threat detected.
pattern:002	Pattern rediscovered.
profile:001	Pattern discovery run started.
profile:002	Pattern discovery run finished.

Query Viewers

Audit Events for Query Viewer Category

Device Event Class ID	Audit Event Description
queryviewer:100	Base query used by the query viewer succeeded.
queryviewer:101	Base query used by the query viewer failed.
queryviewer:102	Base query used by the query viewer has started.

Reports

Audit Events for Report Category

Device Event Class ID	Audit Event Description
report:100	Generated a new archived-report configuration resource.
report:101	Failed to generate a new archived-report configuration resource.
report:102	Generated a new delta archived-report configuration resource.
report:103	Report generation was canceled.
report:104	Report generation started.
report:105	Report generate process halted because the report was empty.
report:106	Report was successfully sent through email.
report:107	Failed to send report through email.

Repository Audit Events

Audit Events for Repository

Device Event Class ID	Audit Event Message
repo:100	Repository is up
repo:101	Repository is down

Resource Quota

Audit Events for Resource Quota Category

Device Event Class ID	Audit Event Description
quota:100	Resource usage has fallen below the fixed-quota level.
quota:101	Resource usage has exceeded the fixed-quota level.
quota:102	Asset autocreation has exceeded a fixed quota.
quota:103	Asset autocreation is proceeding too rapidly.

Rule Actions

Audit Events for Rule Actions Category

Device Event Class ID	Audit Event Description
rule:300	For rule actions that do not have specific DEC IDs assigned.
rule:302	Set Event Attribute action.
rule:303	Send to Notifier action.
rule:304	Execute Command action.
rule:305	Export... action.
rule:306	Create New Case action.
rule:307	Add to Case action.
rule:308	Create New Case action failed.
rule:309	Add to Case action failed.
rule:310	Add to Active List action.
rule:312	Remove from Active List action.
rule:313	Run SmartConnector (agent) command.
rule:315	AddAssetCategory.
rule:316	RemoveAssetCategory.

Rule Activations

Audit Events for Rule Activations Category

Device Event Class ID	Audit Event Description
rule:700	Rule has been deactivated.
rule:701	Rule has been deactivated because it is unsafe. There was excessive recursion or event matching.
rule:702	Rule has been activated.
rule:703	Unsafe rule activation.

Target User Name or Target User ID in the Audit Event Affected by Who Triggers Activation or Deactivation

Whether a rule was disabled or enabled by a:

- User (compact mode and distributed mode)
- System (compact mode and distributed mode)
- Correlator or aggregator (distributed mode only)

affects the target user name or target user ID shown in the resulting audit event, as shown in this table:

Who enables or disables the rule?	Compact Mode		Distributed Mode	
	Target User Name	Target User ID	Target User Name	Target User ID
User	login user name	Target User ID data	login user name	Target User ID data
System	Empty	Empty	Empty	Empty
Correlator or aggregator	Not applicable for compact mode.	Not applicable for compact mode.	arcsightclusteruser	Target User ID data

Rule Firings

Audit Events for Rule Firings Category

Device Event Class ID	Audit Event Description
rule:101	Rule fired OnEveryEvent.
rule:102	Rule fired OnFirstEvent.
rule:103	Rule fired OnSubsequentEvents.
rule:104	Rule fired OnEveryThreshold.
rule:105	Rule fired OnFirstThreshold.
rule:106	Rule fired OnSubsequentThresholds.
rule:107	Rule fired OnTimeWindowExpiration.
rule:108	Rule fired OnTimeUnit.

Rule Warnings

Device Event Class ID	Audit Event Description
rule:501	Rule is firing on events generated by itself (infinite loop)
rule:601	If the max rule chain is exceeded, an audit event with a rate limit of every 30 seconds will be sent with the name: "Exceeded max rule chain <maxRuleChain> for the rule <ruleName>" Device Event Category: /Rule/Error/RuleChain

Rules Scheduled

Device Event Class ID	Audit Event Description
rule:801	Scheduled rule started.
rule:802	Scheduled rule finished.

Scheduler Execution

Device Event Class ID	Audit Event Description
scheduler:200	A task has been executed.
scheduler:201	A task failed to execute.

Scheduler Scheduling Tasks

Audit Events for Scheduler Scheduling Tasks Category

Device Event Class ID	Audit Event Description
scheduler:300	A new task has been scheduled.
scheduler:301	A new task could not be scheduled.
scheduler:302	Enabled a task.
scheduler:303	Could not enable a task.
scheduler:304	Deleted a task.
scheduler:305	Failed to delete a task.
scheduler:306	Disable a task.
scheduler:307	Could not disable a task.

Scheduler Skip

Device Event Class ID	Audit Event Description
scheduler:100	The task scheduler skipped a scheduled task execution because the scheduler was not allowed to run.
scheduler:101	The task scheduler skipped a scheduled task invocation because the last invocation of the task is still executing.

Session Lists

Audit Events for Session List Category

Device Event Class ID	Audit Event Description
sessionlist:101	An entry was added to a session list.
sessionlist:102	An entry was removed from a session list.
sessionlist:103	A session list entry was updated.
sessionlist:104	An entry in a session list was auto-terminated as the session expired.
sessionlist:105	A session list entry was evicted.
sessionlist:201	A session list partition was added.
sessionlist:202	A session list partition was dropped.
sessionlist:203	A session list Partition add failed.
sessionlist:204	A session list Partition drop failed.
sessionlist:301	During lookup on a session list value, the value was not available in Manager memory, and the lookup was not performed on the database. This can occur if too many session list lookups are performed against the database. Typically, the Manager generates one audit event for any number of dropped lookups in a time period, instead of one per dropped lookup.

Transformation Hub

Audit events are generated every five minutes.

Transformation Hub Audit Events

Device Event Class ID	Audit Event Description
thub:100	Connection to Transformation Hub is up
thub:101	Connection to Transformation Hub is down
thub:102	Number of messages remaining in Transformation Hub
thub:103	Number of events forwarded from Transformation Hub to ESM

Trends

Audit Events for Trend Category

Device Event Class ID	Audit Event Description
trend:100	Trend run started.
trend:101	Trend run success.
trend:102	Trend run failure.
trend:201	Trend scavenge success.
trend:202	Trend scavenge failure.
trend:401	Trend enabled.
trend:402	Trend disabled.
trend:501	Trend task started.
trend:502	Trend task ended.
trend:601	Trend was automatically deactivated because of too many failures.
trend:701	Trend successfully updated an active list. You can add an action to a trend to send columns (fields) in trend results to a <i>fields-based</i> active list.

Trend Partitions

Audit Events for Trend Partitions Category

Device Event Class ID	Audit Event Description
trend:301	Trend partition added.
trend:302	Trend partition dropped.
trend:303	Trend partition add failed.
trend:304	Trend partition drop failed.

User Login

Audit Events for User Logins

Device Event Class ID	Audit Event Description
authentication:100	Successful client login.
authentication:101	Failed client login.
authentication:102	Client logout.
authentication:103	Client timed out due to inactivity.
authentication:104	Too many client login failures occurred within a time period. Note: After the third login failure, future logins are prevented. The next time this user logs in, the generated event is authentication:101, Failed client login with the reason, User disabled.

User Management

Audit Events for User Management Category

Device Event Class ID	Audit Event Description
user:100	A user account was deleted.
user:101	A user account was updated. This audit event is generated when an existing user account is modified or a new user is inserted.
user:102	A user account was added. When a new user account is inserted, two audit events are generated: this User Inserted event, and a User Update event (user:101).

Also see the section “Group Management,” earlier in the Audit Events section, which reflects adds, deletes, and updates of groups, including user groups.

Base Queries

A primary attribute of any query viewer, report, and trend is the SQL query it references and uses. If you create the query viewer, report, or trend yourself, you define this as part of the initial attributes by browsing to and choosing a query from the Reports/Queries tree.

For information on creating queries, see ["Building a Query" on page 320](#).

Batching

Batching is a mode in which ArcSight SmartConnectors receive or send collections of events at one time rather than immediately after each occurrence.

Cases

Cases are entries in an event-tracking system used to track, investigate, and resolve suspicious events in a workflow-type environment. When suspicious events occur, cases are created and assigned to users, who then investigate and resolve them based on enterprise policies and practices.

ArcSight has two ways to create and handle cases. First, it has its own complete case-management system. You can use this system to create new cases and assign them to specific groups and users who are notified and receive the cases and relevant data and information associated with the case. Those users can then act on the assigned cases, specifying resolution or other actions taken on the case, which gets reported back and recorded in the ongoing or final resolution of a case.

In addition to using the built-in case management system that ArcSight provides, you can also integrate it with other external case management systems such as Remedy. In that situation, adding new cases exports event information and bring up forms of the external case management system for you to create and assign new cases. The integration with external case management system can also be customized so that case resolution is reported back and recorded.

Case attachments enable you to attach files to any case you are able to edit, for example log files. You are also able to delete cases and attachments; if you delete a case, it deletes the attachment. You can add a file to a case, making it public or private. Private means that the attachment is never shared with other cases; Public means that everyone has access to the latest edited version of that file. Sharing attachments makes it possible to share files that are common among many cases, for example as with a non-disclosure agreement.

For complete information on working with cases, see ["Case Management and Queries" on page 624](#).

Case Groups

Cases are organized into these groups:

Case Group	Description
<User Name>'s Cases	Cases assigned to the user ID.
Shared Cases	Cases that the logged-in user has permission to access.
Public Cases	Cases to which all users have read permission.
Unassigned	Cases that are not assigned to any user.

If you have Administrator access, you also have a group named All Cases that contains all user case groups and their cases.

Categories

ArcSight uses categories and a set of supporting attributes to distinguish the events reported by SmartConnectors or generated internally by Managers. You see these under the Category heading in tools such as the [Common Conditions Editor \(CCE\)](#), [Rules Editor](#), or [Event Inspector](#). This ability to recognize more detailed and specific event conditions increases your analytic and reactive options.

These categories and attributes are designated by ArcSight, based on the information offered to SmartConnectors by sensors. Keep in mind that the applicability of a category always depends on the actual configuration of the environment.

Category Groupings

Category	Description
Object Category	Events are always about a certain object. An object can be an application, the operating system, the database, a file, or the memory of a server. The object is assumed to be the targeted object being accessed, altered, and so on.
Behavior Category	Events not only refer to certain objects, but there is generally an action or a behavior associated with an event. What is being done to an object? Behaviors include access, execution, or modification, and so on.
Outcome Category	After the object is identified and the behavior determined, outcome is the next step. The outcome can be a success, a failure or an attempt. An attempt really indicates that something was neither a success nor a failure and the outcome is not clear or there is no statement that could be made about the outcome.

Category Groupings, continued

Category	Description
Device Group Category	Many security devices serve a multitude of purposes in one product. For example, intrusion prevention systems generate events associated with their firewall capabilities, as well as their intrusion detection capabilities. To be able to distinguish between these types of events, we introduced a dimension called deviceGroup . This dimension allows us to query, for example, for all of the firewall-type events as opposed to all of the events generated by a firewall. The distinction being that the former query also returns all the firewall messages, for example, in the operating system logs (for example, ip tables).
Significance Category	We need to know the significance of an event. We need the capability, for example, to separate normal events from hostile events. We also need to know whether certain activity reported by the device impacts the availability, confidentiality, or integrity of our systems. All this information is captured in the significance category.
Technique Category	Frequently in a security context, we would like to obtain information about the type of events with respect to a security domain. Is an event talking about a denial of service, a brute force attack, IDS evasions, exploits of vulnerabilities, and so forth.

Object Category

Object Category Details			
Host	Any end-system on the network, such as a PDA, a Windows computer, or a Linux computer.		
	Operating System	The system software that controls execution of computer programs and access to resources on a host.	
	Application	A software program that is not an integral part of the operating system.	
		Service	An application that normally executes at operating system startup. A service often accepts network connections.
		Database	A database application.
		Backdoor	An application, visible on a host, that listens for network connections and can give a non-authorized user control over that host.
		DoS Client	A host that is displaying an application that can participate in a (possibly distributed) denial-of-service attack.
		Peer to Peer	An application that listens for, and establishes network connections to, other installations of the same application (for example, Kazaa, Morpheus, Napster).
		Virus	A host that is displaying a replicating infection of a file that also executes other behaviors on the infected host.
		Worm	A host that is displaying a self-replicating program that spreads itself automatically over the network from one computer to the next.

Object Category Details			
	Resource	An operating system resource that is characteristically limited in its supply.	
		File	A long-term storage mechanism (for example, files, directories, hard disks, and so forth).
		Process	A single executable module that runs concurrently with other executable modules.
		Interface	An interface to the network.
		Interface Tunnel	Packaging a lower network protocol layer within a higher layer (for example, IPSec Tunnel, HTTP tunneling).
		Registry	The central configuration repository for the operating system and the applications. Application-specific information is not stored here.
		CPU	Events directed at this object relate to consumption or use of the overall processing power of the host.
		Memory	Events directed at this object relate to consumption or use of the overall memory of the host.
Network			Events that cannot be clearly associated with a host's subitem. Events that involve transport, or many hosts on the same subnet.
	Routing		Routing related events such as BGP.
	Switching		Switching related events such as VLANs.
Actor	User		A single human identity.
	Group		A named collection of users, such as an employee division or social group.
Vector			The replication path for a section of malicious code.
	Virus		A replicating infection of a file that also executes other behaviors on the infected host.
	Worm		A self-replicating program that automatically spreads itself across the network, from one computer to the next.
	Backdoor		An application that listens for network connections and can give a non-authorized user control over that host.
	DoS Client		An application that participates in a (possibly distributed) denial-of-service attack.

Behavior Category

Behavior Category Details		
Access	Refers to accessing objects, as in reading.	
	Start	The start of an ongoing access, such as login.
	Stop	The end of an ongoing access, such as logging out.
Authentication	Actions that support authentication.	
	Add	Adding new authentication credentials.
	Delete	Deleting authentication credentials.
	Modify	Modifying authentication credentials.
	Verify	Credential verification, such as when logins occur.
Authorization	Authorization-related actions.	
	Add	Adding a privilege for the associated object (for example, a user).
	Delete	Removing a privilege for the associated object (for example, a user).
	Modify	Modifying the existing privileges for the associated user or entity.
	Verify	An authorization check, such as a privilege check.
Communicate	Transactions that occur over the wire.	
	Query	Communicating a request to a service.
	Response	Communicating a response to a request, from a service.
Create	Seeks to create resources, install applications or services, or otherwise cause a new instance of an object.	
Delete	The reverse of creation events. Includes uninstalling applications, services, or similar activity.	
Execute	Involves loading or executing code, booting or shutting systems down, and similar activity.	
	Start	The beginning of execution of an application or service. This event is clearly distinguished from a lone "Execute" attribute.
	Stop	The termination of execution of an application or service. This event is clearly distinguished from a lone "Execute" attribute.
	Query	A query sent to a specific entity - but not over the network (for example, as when generating a report).
	Response	The answer returned by an Execute/Query. For example, a report delivered back from an application, or status messages from applications.
Modify	Involves changing some aspect of an object.	

Behavior Category Details		
	Content	Changing the object's content, such as writing to or deleting from a file or database.
	Attribute	Changing some attribute of an object, such as a file name, modification date, or create date.
	Configuration	Changing an object's configuration. For example, application, operating system, or registry changes.
Substitute	Replacing files, upgrading software, or service or host failovers.	
Found	Noticing an object or its state.	
	Vulnerable	An exploitable state that is characteristic of a particular hardware or software release.
	Misconfigured	An exploitable state caused by a weak configuration or similar mishandling.
	Insecure	An exploitable state that arises from poor management or implementation. For example, weak authentication, weak passwords, passwords passed in the clear, default passwords, or simplistically named accounts.
	Exhausted	The targeted object was found to be exhausted (for example, not enough file descriptors available).

Outcome Category

These attributes indicate the probable success or failure of the specified event, within an overall context. For example, the outcome of an event such as an operation failed error message can be reported as a /Success given that the operation can be presumed to have actually caused a failure. Another example would be an event that identifies a Code Red infection: on a host running Linux the outcome would be /Failure (Code Red is Windows-only) while the same event directed at a host with an unknown OS would be reported as an /Attempt.

Outcome Category Details	
Attempt	The event occurred but its success or failure cannot be determined.
Failure	The event can be reasonable presumed to have failed.
Success	The event can be reasonable presumed to have succeeded.

Device Group Category

Device Group Category Details			
Application	An application program		
Assessment Tool	A network- or host-based scanner that monitors issues such as vulnerability, configurations, and ports		
Security Information Manager	A security-event processing correlation engine (the Manager). This "device" deals only in correlated events.		
Firewall	A firewall		
IDS	An intrusion-detection system. Includes:		
	Network	A network-based intrusion-detection system	
	Host	A host-based intrusion-detection system. Includes:	
		Antivirus	An anti-virus scanner
		File Integrity	A file-integrity scanner
Identity Management	Identity management		
Operating System	An operating system		
Network Equipment	Network equipment. Includes:		
	Router	A network device with routing (layer 3) capabilities	
	Switches	A network device with switching (layer 2) capabilities	
VPN	A virtual private network		

Significance Category

Significance Category Details		
Hostile	A malicious event has happened or is happening.	
Informational	Events considered worthy of inspection; for example, those produced by polling. Includes:	
	Error	An execution problem.
	Warning	A possible problem.
	Alert	A situational problem that requires immediate attention.
Normal	Ordinary or expected activity that is significant only for historical analysis purposes.	
Recon	Relates to scans and other reconnaissance activity.	
Suspicious	A potentially malicious event occurred.	

Technique Category

Technique Category Details			
Traffic	An anomaly in the network traffic, such as non-RFC compliance.		
	Network Layer	Anomalies related to IP, ICMP, and other network-layer protocols. Includes:	
		IP Fragments	Fragmented IP packets.
		Man in the Middle	A man-in-the-middle attack.
		Spoof	Spoofing a source or destination IP address.
		Flow	A problem in network-layer communication logic, such as an out-of-order IP fragment.
	Transport Layer	Anomalies related to TCP, UDP, SSL, and other transport-layer protocols. Includes:	
		Hijack	Hijacking a connection.
		Spoof	Spoofing a transport layer property (for example, a TCP port number, or an SSL entity).
		Flow	A problem in TCP connections or flows, such as a SYNACK without SYN, a sequence number mismatch, or time exceeded.
	Application Layer	Application-layer anomalies. Includes:	
		Flow	A peer does not follow the order of commands.
		Syntax Error	A syntax error in an application-layer command.
		Unsupported Command	A command which does not exist or is not supported.
		Man in the Middle	A man-in-the-middle attack on the application layer.
Exploit	Vulnerability	Exploiting a vulnerability (for example, a buffer overflow, code injection, or format string).	
	Weak Configuration	Exploitation of a weak configuration. This is something that could be remedied easily by changing the configuration of the service (for example, weak passwords, default passwords, insecure software versions, or open SMTP relays).	

Technique Category Details		
	Privilege Escalation	A user identity has received an increase in its user privileges.
	Directory Transversal	A user identity is attempting to browse or methodically review directories for which it may not have appropriate privileges.
Brute Force	Brute-force attacks. Includes:	
	Login	Continued trials for logins.
	URL Guessing	Continued trials for URLs to access information or scripts.
Redirection	Redirecting an entity. Includes:	
	ICMP	ICMP redirects.
	DNS	Unauthorized DNS changes.
	Routing Protocols	Attacks aimed at routing protocols (for example, BGP, RIP, OSPF).
	IP	Redirection using the IP protocol (for example, source routing).
	Application	Redirection attacks on the application layer (for example, cross-site scripting, mail routing, or JavaScript spoofing).
Code Execution	Either the execution or transmission of executable code, or the transmission of a distinctive response from executed code. Includes:	
	Trojan	The code in question is concealed within other code that serves as a Trojan Horse. In other words, it appears to be one thing (that is safe) but is really another (that is unsafe).
	Application Command	The code in question is intended to invoke an application command.
	Shell Command	The code in question is intended to be executed in a shell.
	Worm	Code associated with a worm.
	Virus	Code associated with a virus.
Scan	Any type of scanning. A network, host, application, or operating system scan can be identified through the specified object. Includes:	
	Port	Multiple ports are scanned.
	Service	A service is scanned (for example, DDoS client discovery, backdoors, RPC services, or scans for a specific application such as NMB).
	Host	Scanning for hosts on a network.
	IP Protocol	A search for responding protocols. Note that TCP and UDP are not the only transport protocols available.
	Vulnerability	A scan for vulnerabilities.
DoS	A denial of service attack is in progress.	

Technique Category Details		
Information Leak	Information leaking out of its intended environment (for example, mail messages leaking out, system file access, FTP data access, or web document access). Includes:	
	Covert Channel	Leakage was detected from a covert channel, such as Loki.
Policy	Policy-related violations such as pornographic web site access. Includes:	
	Breach	A policy-related security breach occurred.
	Compliant	A policy-compliant event occurred.

Asset Categories

Asset categories are system resources that describe properties of an asset, such as the operating system running on it, key applications it hosts, its role within the enterprise, and any other properties you want to consider when evaluating threats or behaviors associated with this asset.

There is more information about asset categories in *ESM 101*. See “Asset Categories” in the main topic, “The Network Model.”

The Asset Categories subtab in the Navigator provides options to organize assets into groups based on categories. See ["Managing Asset Categories" on page 171](#).

Event Categories

Events received from ArcSight-supported devices are automatically **categorized** (appended with classification information) by ArcSight SmartConnectors based on the ArcSight event categorization taxonomy. Event Categories are used to classify events based on criteria such as object type, behavior, outcome, technique, device group, and significance. This additional information about an event (together with normalization and other filtering) helps to identify the significance events from different devices and vendors based on a consistent model.

For information about categorization of events by SmartConnectors go to the Protect 724 Web site and download the document entitled *ArcSight Categorization White Paper*.

For information on creating custom event categories, see ["Event Categorization" on page 997](#).

Collaboration

The ArcSight Console's collaboration capability is a collection of features that make it possible for analyst teams to select certain security events for further on-going investigation.

Investigation involves a workflow-style process of information collection that leads through a series of analysis stages to a final disposition.

In the ArcSight Console, you locate events for analysis through the active channels in the Viewer panel. You use the Annotate Events dialog box or the Event Inspector to annotate an event, or collection of events, and set it up for follow-on analysis. Once you have placed the event or collection in the collaboration "pipeline" by assigning it a disposition stage (such as **Initial**), you or other analysts manage it through to resolution using the assigned stages as filter arguments.

ArcSight provides a set of default collaboration stages, but your enterprise may well use others created specifically for your workflow needs. Owners, disposition stages, comments, and other factors change as an event's handling progresses. The collaboration cycle usually ends when someone marks the event's **Stage** field as **Closed** (or the equivalent).

Compare collaborative annotation to cases, which are a more formal way to track sets of events that are under investigation.

Default Collaboration Stages

Stage	Meaning
Queued	The event has not yet been inspected.
Initial	The event has been inspected.
Follow-up	The event is under investigation.
Final	The investigation has concluded.
Closed	The investigation is closed.
Monitoring	The event is being watched, especially in regard to a reoccurrence in support of a pattern.
Rule Created	The event has been used to create a rule that assists in monitoring for a reoccurrence, especially in regard to patterns, and potentially to respond in some way such as with a notification.
Flagged as Similar	The event is similar to one already under investigation.

See ["Collaborating on Events \(Event Annotation\)" on page 295](#) for descriptions of the tasks involved.

Common Conditions Editor (CCE)

The ArcSight Console has a Boolean logic editor, which is also referred to as the Common Conditions Editor (CCE). If the criteria are met, the evaluation returns a Boolean true or false. All conditions constructed by the CCE are expressions that consist of a value or variable on the left, an operator (NOT, AND, OR), and a value on the right, for defining the conditions you use

to help analyze resources such as filters, rules, and reports. This topic is your reference for using the CCE, wherever it appears.

Topics:

- ["Editor Features" below](#)
- ["Condition Tree Command Buttons" on page 890](#)
- ["Condition Tree Context Menu Commands" on page 892](#)
- ["Adding Conditions" on page 895](#)
- ["Using Field Sets" on page 899](#)
- ["Adding or Removing Global Variables Using the CCE" on page 901](#)
- ["Testing for Zone Relevance" on page 902](#)

See also:

- ["Logical Operators" on page 1006](#) for related information in creating conditions.
- ["Filtering Events" on page 306](#), especially subtopics on [Creating or Editing a Filter](#) and [Debugging Filters to Match Events](#).

Editor Features

The CCE has two tabs; **Edit** and **Summary**. In the **Edit** tab, logical operators are represented in a tree form.

In the **Summary** tab, conditions are presented in an easily readable, summary view. Resource references in the **Summary** tab are hyperlinked. From the Summary tab, click a resource link to open its definition in a resource editor in the Inspect/Edit panel.

Conditions are editable only on the **Edit** tab. Wherever the CCE appears, you use these features to build or change conditional expressions.


- The condition tree shows the complete set of expressions you are building or changing.
- The root of the tree indicates whether the expression concerns [Filters \(Filter By\)](#), [Correlation \(Correlate\)](#), or [Reports \(Report On\)](#), as you see in the Filters Editor, Rules Editor, or Report Editor, respectively.
- From the root, there are branches for one or more events. For each event branch, there are sub-branches for one or more condition statements.
- To add an event for a rule, select the root and click the **New Event Definition** button (see below) or right-click the root and choose the same command. Note that only rules can add events because filters and reports do not need additional events for correlation.
- To act on a specific event or [Conditional Statements](#), select it in the tree. Once selected, you can use several features to modify it, as described here and below.

- Use the new event, [Logical Operators](#), and resource selector buttons above the tree to add events, operators, or resource-based constraints to condition statements, if applicable.
- Use the right-click menus that are available for any selected branch of the condition tree to choose commands that are applicable to that statement in that context.
- When you use the right-click **Edit** command to edit a selected statement directly in the tree (rather than through the event fields table), you can use the Enter key to update the condition without having to click **Apply** or **OK**.
- Do use single- or multiple-selection copying and pasting of statements for efficiency. You can use the right-click menu commands or **Ctrl+C** for copy, **Ctrl+X** for cut, and **Ctrl+V** for paste.
- Use the [Field Sets](#) selector to choose an appropriate group of event fields when an event-related statement is selected in the condition tree.
- To **undo/redo** an action, right-click in the Edit panel and choose either **Undo** or **Redo**, depending upon the action you want to use. For example, if you decide to delete a node, a message asks you to confirm. If you want to undo this delete, right-click in the Edit panel and choose **Undo Delete**. (You can also use the standard keyboard commands **Ctrl+Z** for undo and **Ctrl+Y** for redo.)
- To **Search** for a resource, simply click in the field column (on the left side of the list) and start typing. A Search popup is displayed when you start typing, and shows the term as you type it. The search is "predictive" in that it navigates to and select matching fields as you type. Click **Enter** to select this resource. For details see ["Searching for Fields in Event Inspector, Resource Editors, or CCE" on page 70](#).




Note: Both tabs provide syntax and error highlighting. As an example of error highlighting, if a condition uses resources that are later removed, references to the missing resources are highlighted as errors in the condition statements in the CCE.

In the **Conditions "Edit"** tab, logical operators are represented in a tree form. Use this tab to define and edit conditional statements.



In the **Conditions "Summary"** tab, conditions are shown in an easily readable, summary view. In this view, resource references are hyperlinked. Click on the links to go to a resource definition. Use this tab to review condition statements. You cannot edit the Condition on this tab.



Condition Tree Command Buttons

All of the following options are available from buttons at the top of the Conditions Editor and also from right-click menu options. The exception is the "In Case" Condition which is only available from a right-click menu option.

Condition Tree Command Buttons

Button	Name	Use
	New Event Definition	Insert a new condition tree in the editor.
	AND	Insert an AND condition.
	OR	Insert an OR condition.
	NOT	Insert a NOT condition.
	Filters	Matches Filter condition. This resource-based command browses the Filters tree of the Navigator panel. Note that this operator applies only to rules.
	Assets	Assets condition. This resource-based command browses the Assets tree of the Navigator panel. Note that this operator applies only to rules.
	Vulnerabilities	Vulnerabilities condition. This resource-based command browses the Vulnerabilities tree of the Navigator panel. Note that this operator applies only to rules.
	Active Lists	Active Lists condition. This resource-based command browses the Active Lists tree of the Navigator panel. Note that this operator applies only to rules.
	Joins	Matching Event condition. Applies when there are two or more event conditions.
	Vulnerabilities	Has Vulnerability condition. This resource-based command browses the Vulnerabilities tree in the Navigator panel.
	Active List	InActiveList condition. This command browses the Navigator panel's Active Lists tree, and operates on items in the event and actor schemas. It is used to map a field or a global variable in the event schema to a corresponding field in an active list. It does not evaluate items in other non-event schemas (such as cases or assets). The InActiveList operator option evaluates single-value attributes and multi-value attributes. The field you map could return multiple values. In the case of multi-value attributes, if any one value matches, the condition evaluates to true. Consider this scenario for multi-value attributes: An active list keeps track of actor roles where role values can be one of Normal, Restricted, or Privileged. You can test if an actor has one of these roles. If your list has a field called RoleName, you map the actor's RoleName attribute to this field. Keep in mind that an actor's RoleName attribute is multi-valued because an actor can have multiple roles. Through the InActiveList condition, you can have a query that checks if one of the actor's roles is Privileged. A condition that tests for whether all or any values in a list match is only available to specify on queries and on in-memory operations such as rules, filters, data monitors. Note: The InActiveList condition in lightweight rules does not support lists with multi-mapped values.
	Joins	Inserts a Join or Matching Event condition. Note: This option applies only to Rules. See "Creating Matching or Join Conditions" on page 519.

Condition Tree Context Menu Commands

Condition Tree Context Menu Commands

Command	Description	Applies To
New Condition	Add a new condition statement below the selected element. Type the statement directly in the tree or choose a field from the pop-up menu.	operator, event field
New Logical Operator	Add a new logical operator to the selected element. See "Logical Operators" on page 1006 .	Event alias, operator, event field
New Constant Condition	Add a Boolean (True/False) AND operator to the selected branch.	operator
New "Matches Filter" Condition	Use the Filter selector to identify a particular filter as a matching argument for a condition. See also "Creating Matching or Join Conditions" on page 519 .	operator, event field
New "Assets" Condition	Use the Assets selector to identify an asset or group as the argument for a condition. See also "Adding Asset Conditions" on page 516 .	operator, event field
New "Has Vulnerability" Condition	Use the Vulnerability selector to identify a vulnerability as the argument for a condition. See also "Adding Vulnerability Conditions" on page 516 .	operator, event field

Condition Tree Context Menu Commands, continued

Command	Description	Applies To
New "InActiveList" Condition	Use the Active List selector to identify a particular active list that contains the argument for a condition. It is used to map a field or a global variable in the event schema to a corresponding field in an active list. It does not evaluate items in other non-event schemas (such as cases or assets). When the InActiveList condition is used to compare values in two lists, an additional option is shown where you can specify whether All values in list field must match. If this option is checked, the Active List condition evaluates to true only if all values in both lists match. If it is not selected, the condition evaluates to true if any field is in both lists. Note: The InActiveList operator option evaluates single-value attributes and multi-value attributes. The field you map could return multiple values. In the case of multi-value attributes, if any one value matches, the condition evaluates to true. Consider this scenario for multi-value attributes: An active list keeps track of actor roles where role values can be one of Normal, Restricted, or Privileged. You can test if an actor has one of these roles. If your list has a field called RoleName, you map the actor's RoleName attribute to this field. Keep in mind that an actor's RoleName attribute is multi-valued because an actor can have multiple roles. Through the InActiveList condition, you can have a query that checks if one of the actor's roles is Privileged. A condition that tests for whether all or any values in a list match is only available to specify on queries and on in-memory operations such as rules, filters, and data monitors. See also "Adding Active List (InActiveList) Conditions" on page 517.	operator, event field
New "InCase" Condition	Use the Cases resource tree to identify a particular case as the argument for a condition. For events, the condition checks if the event is part of a case's details. Note: The channel that uses a filter with the InCase condition is set to evaluate only once, therefore making the channel static. In this situation, if a rule subsequently adds an event to the case as you are viewing the case's events on the channel, the channel is not updated until you close, then re-open (refresh) the channel. See also: • "Applying Rule Actions on Cases" on page 546 • "Working with Events in Cases" on page 640	operator, event field
New Event Definition	Create and name a new event alias to add to the root. Note: This option applies only to Rules.	root
Change Operator	Change the rule operator to And , Or , or Not .	operator

Condition Tree Context Menu Commands, continued

Command	Description	Applies To
Set Global Expiration Time	For rules, set the amount of time that qualifying events for all aliases are retained in memory for evaluation, based on Manager receipt-time. Setting an alias expiration overrides a global expiration, if present. See "Specifying Rule Thresholds and Aggregation" on page 528 for more information. Note: This option applies only to Rules.	root
Align Nodes	When selected, shows the hierarchical structure of event conditions. Note: This option applies only to Rules.	root
Edit	Open a text box in which to change the selected element.	operator, event field
Undo	Undo an action.	all actions
Redo	Redo an action.	all actions
Cut	Cut the selected elements of the condition tree to the Clipboard.	root, event alias, operator, asset, event field
Copy	Copy the selected elements of the condition tree to the Clipboard.	root, event alias, operator, asset, event field
Paste	Paste the conditional element currently on the Clipboard to the end of the selected element in the tree.	root, event alias, operator, asset, event field
Delete	Delete the selected elements of the condition tree.	event alias, operator, asset, event field
Set Alias Expiration Time	For rules, set the amount of time that a qualifying event for this alias (only) is retained in memory for evaluation, based on Manager receipt-time. See "Specifying Rule Thresholds and Aggregation" on page 528 for more information. Note: This option applies only to Rules.	event alias

Condition Tree Context Menu Commands, continued

Command	Description	Applies To
Consume After Match	Use the event only once to fire a rule. Thereafter, additional joins with other event conditions are not performed within the rule's time window. This setting is used to reduce the number of correlation alerts. By default, this setting is off. If disabled, an event matching a rule's event condition alias stays in working memory and continues to combine with events that match other aliases, until the event itself expires within the time window. Note: This option applies only to Rules.	joins
Negated	For rules, a way to monitor the non-occurrence of an event. See "Negating Event Conditions" on page 522 for details on how to trigger rule actions based on negated events. Setting an event condition to Negated requires you to enter a timeout value. If the negated event is not sent from the device within this period, the rule is triggered. Note: This option becomes available if the rule has two or more event conditions.	event condition alias
Set Matching Time	Sets the maximum time difference between the partially-matched aliases. Note: This option applies only to Rules.	matching event operator
Print Conditions and Tree Summary	Prints the condition definition as shown on the Edit tab and the Summary statement. Selecting this menu option brings up a Print Preview dialog where you can view what will print, and set printer options.	event alias, operator, asset
Help	Open the online Help system for information about the type of resource being edited.	root, event alias, operator, asset, event field

Adding Conditions

When adding conditions, decide how the new condition ties to existing conditions. If AND is used, the new condition has to occur in addition to existing conditions. If OR is used, the new condition or any existing conditions have to occur. If NOT is used, all but the new condition has to occur.

You use the AND, OR, and NOT operators to define relationships between condition statements. When you use AND, the new condition must occur in addition to the selected condition. Using OR means either the selected or new condition must occur. Using NOT means all but the new condition must occur.



Tip: Multiple assets and asset categories added to a single asset condition are always OR'ed together (not AND'ed).

For example, create a new rule, click the **Conditions** tab in the Rule Editor, select **Assets**, and add some asset categories to the condition. (To do this, select them on the Asset Categories tab at the bottom of the Editor and click **Apply**.)



Click the **Summary** tab to view the detail of the Boolean logic. This shows that the assets are OR'ed together.



If you want to AND an asset condition to other conditions, go back to the **Edit** tab, select the event definition again, and add other conditions based on the fields shown in the lower half of the editor.

To add more condition statements, right-click an existing statement and choose **New Logical Operator**, then **And**, **Or**, or **Not**, or click a logical operator or resource-selection button. Then, create the new condition statement.

Event-definition and Join conditions are allowed only with rules to include separate events or aliases, or correlation of these separate events respectively.

In the data field table, scroll to a data field in the Name column to create a condition statement.

Data fields provide event details from all devices deployed throughout your enterprise. Event details from these devices are normalized into common data fields and stored in the database to allow investigative and analytical comparison of all incoming events. See ["Data Fields" on page 910](#) and ["Timestamp Variables" on page 1060](#) for more information.

The data field table displays a **Name**, **Operator**, and **Condition** column. These three columns are combined to create <data field> <logic operator> <data field value> condition statements. For example, if monitoring a Cisco Router, you could define a condition statement

to specify Device Product = Cisco Router: Device Product as the data field, equals (=) as the logic operator, and Cisco Router as the data field value.

See also:

- ["Search Box to Find Fields in the List" below](#)
- ["Field Comparisons with Variable or Static Values" on the next page](#)
- ["Creating Matching or Join Conditions" on page 519](#) (if you are creating rules)

Search Box to Find Fields in the List



Tip: Search Shortcuts

- Type part of the field name you want to find (for example, Name) in the Search box.
- Use the up/down arrow keys to jump to each instance of **Name** in the available fields.
- When you find the field name you want, press Return to add it to the condition statement
- Ctrl+F re-displays the Search box back if it's hidden.


Search Box Example

Search Shortcuts:
Type part of the field name you want to find (e.g., Name) in the Search box.

Use the up/down arrow keys to jump to each instance of "Name" in the available fields.


When you find the field name you want, hit Return to add it to the selected query structure sections (SELECT, GROUP BY, or ORDER BY)

Ctrl+F gets the Search box back in display if it's hidden



Field Comparisons with Variable or Static Values

For any field comparison, a drop-down menu of variables is provided for the *right* side of the statement. Or you can type a value here.




The CCE provides a field comparison ability that allows you to compare one field to another field (for example, AttackerHostName = AttackerUserName). This functionality is available on the Console wherever the CCE is available (in [Rules](#), [Reports](#), [Filters](#), and so on). If the fields you are comparing are numeric, the fields can be of different numeric types, for example, a long type compared to a floating point type.

Left-side event attributes can be compared to right-side conditions (represented as variables or static values) using operators like equals (=), is not equal to (!=), is less than or equal to (<=), is greater than or equal to (>=), is less than (<), is greater than (>), and so forth (see ["Logical Operators" on page 1006](#)).

Using Field Sets

The Common Conditions Editor provides access to all available [Field Sets](#) you created. You can specify fields with particular values as part of conditions statements. See also ["Creating a Field Set" on page 601](#).

You can select a particular field set, which limits the fields shown to a subset of all available field sets. If you cannot find a field, click the "Clear field set" button  to clear the field set selection and show the complete list of field sets. This clears the field set selection and shows the complete list of field sets. A common problem is having the common conditions editor (CCE) field display limited to a field set that does not include some fields you want to use in the condition.



For example, suppose you define a condition to look for two matching events; one in which Event Name contains "swipe" and another in which Event Name contains "login". You can set this condition with the "Standard" field set shown above because it includes the Event Name field in the list of available fields from which to choose. But if you wanted to add conditions based on an Event field for "Correlated Event Count" or Threat field for "Model Confidence," you would clear the Field Set and view all fields to get access to these fields.




Tip: Fields shown in italics are *derived* from data in other fields. Derived fields show up in various places on the Console UI including on the Field Set editor, and the Common Conditions Editor (CCE) aggregation tabs (for example, Rules, Filters, and so forth).

Adding or Removing Global Variables Using the CCE


The Common Conditions Editor enables you to define a local variable to apply to the condition statement for that resource, and it also enables you to place [Global Variables](#) in the condition by using the **+/- Global Variables** button (next to the Field Set selector) on the CCE.

To add global variables and make them available for conditional statements in a resource:

1. In the CCE for a given resource, click the **+/-Global Variable** button . The Global Variable Selector displays the Fields resource tree containing your selection of global variables.
2. Select one or more variables you want to add and click **OK**. The variables are added as part of the available fields on the CCE under the Variables group.
3. On the CCE, scroll to the bottom of the available fields. You can use these variables in condition statements for this resource.



To remove one or more global variables from the available fields list in the CCE for a resource:

1. In the CCE for a given resource, click the **+/-Global Variable** button .
2. On the Global Variable Selector dialog, click to de-select one or more variables you want to

remove and click **OK**.

3. The variables are removed from the list of available fields in the CCE.

More information:

- ["Global Variables" on page 609.](#)
- ["Variables" on page 1064.](#)
- ["Velocity Templates" on page 1088.](#)

Testing for Zone Relevance

Events include several [Data Fields](#) that are related to zones (see ["Assets" on page 830](#)). In the Common Conditions Editor you can compare these fields with asset groups or categories, to test whether the field's event does or does not correlate with those asset properties. This comparison is performed by the **InGroup** operator.

For example, if an event's Attacker Zone field value and a Source Asset ID's System Asset Categories' Criticality value correlate, then the InGroup operator would test True. You can apply this outcome in your reports, rules, or filters.



Note: The InGroup operator is inserted automatically when you create zone-asset correlation statements in the Common Condition Editor. There is no button or command to manually insert the operator.

The InGroup operator tests True for specified asset resources and their parents but not for their own peers or their parent's peers.

1. In the Conditions tab of any appropriate editor, set a logical operator for a zone-related field (for example, Destination Zone).
2. In the same field, click the ellipses button (...). In the Select a Zone dialog, enter a prompt for the condition, select the **Parameter** checkbox, then choose a zone from the resource tree.
3. Right-click the new statement in the editor and choose **AND**, then right-click the AND statement and choose **New Assets Condition**.
4. In the Asset resources panel below, choose the Source, Target, or other type of relevant asset ID.
5. For that asset ID type, click the **Assets** or **Asset Categories** tab and select an asset group or category to test with the InGroup operator.
6. Click **Apply** in the Assets resources panel to add the asset group or category to the condition statement, with the embedded Ingroup operator.

Conditional Statements

This table offers sample conditional expressions you can create using various operators, event fields, and data types.

Data Types in Conditional Statements

ArcSight Data Types	Description
Number or Integer, including MAC Address	Using numeric (integer) fields, you can specify operators including =, !=, <, <=, >=, >, and In to specify a numeric comparison expression, for example: CustomNumber1 = 50. To use In, you can specify any number of comma-separated values to match (or equal). Use the above operators for MAC addresses, for example: Attacker Mac Address != <Mac address>
String	Using string fields, you can specify operators including =, !=, and In, Contains, Matches, Starts With, Ends With, and Like to define a string comparison expression. For example: ArcSightCategory StartsWith /Attack or ArcSightCategory = /AttackSuccess
DateTime	Using DateTime fields, you can specify operators including =, !=, Between, In, and On to specify a datetime comparison expression. For example: DetectTime Between 4/1/03 11:30:01AM, 4/1/03 4:30:01PM. You can enter DateTime values directly or click the ellipsis (...) button to select a date from a pop-up calendar or a special date keyword list. Special date keywords you can use are: Now, 1 or 2 hours ago, 1 or 2 days ago, 1 or 2 weeks ago, or a replay start and end time. You can also use special system variables such as: • \$CurrentDateTime: for the date and time the report is run; the system variable is replaced by the current date and time value. • \$CurrentDate: for the date the report is run; the system variable is replaced with the date value, truncating the time of the day to 0, when the report is scheduled or run. You can specify certain date operations with these system variables to add or subtract a number of specified days or hours. For example, you could type: \$CurrentDate - 7d for seven days before the date the report is run, the condition evaluates to a date which is the current date minus seven days, or \$CurrentDateTime - 12h, which evaluates to the current date time minus 12 hours.
IP Address	Using IP address fields, you can specify operators including =, !=, In, InSubnet, and Between to specify an IP comparison expression. For example: TargetAddress = 192.0.2.0. With the In operator, you can also specify a comma-separated list of IP addresses to match. With InSubnet, you can specify an address range. For IP address range formats, see IP Address Ranges. Caution: For the InSubnet operator, do not mix IPv4 and IPv6 addresses within the same IP address range.

These same rules apply to the conditions editor used in defining rules, creating conditional reports, and filters.



Tip: Using variables


You can use all of the dynamic time parameters you see in the Active Channel Editor and elsewhere, such as `$Now` and `$CurrentDateTime`. The same is true for time elements, including s (second), m (minute), d (date), M (month), w (week), and y (year). To use any event data field as a variable, express its displayed name as a one-word, camel case string prefixed with a dollar sign; for example, "Source Address" is `$sourceAddress`. See the complete discussion in the topic ["Variables" on page 1064](#).

Conditions

Conditions are logical expressions (see ["Logical Operators" on page 1006](#)) used to qualify [Events](#) or other grouping of elements. Conditions can be specified in a number of places using a common condition editor; for example, to define rules or filters.

Parameterized Conditions

Some conditions can be parameterized, for example in reports, where the exact value specified for a condition match is provided at the time you are running the report. You do this through the Reports Parameter popup. This lets report authors define default parameter values for the report with the option to override for a specific report run.



This is a convenience for people running the report as it does not require write permissions while running the report, but effectively provides the same flexibility as being able to modify the report. Note that when defining parameters for detect time, you should always include a BETWEEN condition so that the report is limited to a certain time range, and does not scan the entire event table. Otherwise, it can severely impact the Manager information-retrieval performance.

1. You can select the ellipsis (...) button and then select the **Parameter** checkbox to create a parameter prompt for selected data fields of a report. When users run the report, they are first prompted to enter values for these parameters. When specifying a report parameter, you can define the prompt that is displayed to users, as well as specify a default value that is displayed in the prompt field.
2. In the case-sensitive column (**Aa**), select the checkbox if the data field needs to be case-sensitive.
3. In the negate condition column (the "No" symbol), select the checkbox to change conditions to **all but this** statements.

For example, if the condition statement is Device Product = Cisco Router and the negate condition checkbox is selected, all events but those from the Cisco Router generate a correlation event.



Note: You can set some conditions such that the exact value is provided at the time of running the report through a parameter pop-up box. This lets report authors give default parameter values, which can be overridden by the user running the report. This is a very useful convenience for people running the report as it does not require write permissions while running the report but effectively provides the same flexibility as being able to modify the report. Note that when defining parameters for detect time, you should always include a BETWEEN condition so that the report is limited to a certain time range, and does not scan the entire event table. Otherwise it can severely impact the performance of Manager, retrieving information from the database.

4. Click outside the data field row.
The condition statement (<data field> <logic operator> <data field value>) appears as a branch under the logical operator.
5. On the Conditions tab, click **OK**.

These same rules apply to the conditions editor used in creating other resources such as rules and reports.

Content

ArcSight provides preconfigured [Resources](#), also known as *content*, in the form of packages. Content packages are automatically installed to provide ready-to-use resource suites that you can start using immediately to monitor and protect your network. You can develop your own custom content with the editors and tools provided.

Content Packages

ArcSight ships with system content already developed that addresses common security and regulatory use cases. These use cases combine many resources to address multi-faceted issues. You can use system content as is, or modify it with data specific to your network environment. System content is delivered as packages. (See [Packages](#).) System content packages are automatically installed as a part of ArcSight to provide out-of-box resource suites that you can start using immediately to monitor and protect your network. (See also [Resources](#).)

The content packages provided with a standard installation are:

- ArcSight Administration
- ArcSight System

Custom Content

The term "custom content" refers to resources and solutions created by customers using ArcSight software. Examples of custom content are user-configured [Rules](#), [Filters](#), [Active Channels](#), [Queries](#), Trends ([Building Trends](#)), and [Reports](#) designed to address customer-specific scenarios.

SmartConnector Content

ArcSight continuously develops new SmartConnector event categorization mappings, which are often called "content." All existing content is included with major product releases, but it is also possible to get regular content updates to stay completely current.

ArcSight makes available to subscribing customers downloadable packages of new SmartConnector content. These releases occur frequently, generally on a bi-weekly basis. The download files are offered through a special subdirectory on a download site. This directory is visible only to subscribers, who receive a notification e-mail from Customer Support when the files are posted.

The content is packaged in .aup (update) files which may be wrapped in .ZIP files for transmission convenience. You place these .aup files in the Manager's /updates directory, where the Manager automatically finds the content and pushes it to the SmartConnectors. The affected SmartConnectors each register an internal event when the update occurs and can notify you by e-mail through the Manager.

Contact your Customer Support representative to arrange for a content subscription.

CORR-Engine

The CORR-Engine is the central repository for all events. Once an event occurs, its Data Fields such as severity, create time, rules triggered, and so forth are stored in event archives in the CORR-Engine. The CORR-Engine stores all enterprise events in a normalized schema. You can then investigate and analyze the event information. The Manager is the only component that communicates with the CORR-Engine. Also see ["Schema" on page 1043](#).

Correlation

Correlation is logically linking events based on multiple conditions.

See ["Rules Authoring" on page 507](#) and ["Identity Correlation" on page 565](#) to know how correlation is applied.

Correlation Formula

The event correlation data monitor applies covariance and correlation calculations to describe how two variables are related.

Covariance is calculated by the following formula:

$$COV(x,y) = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{n - 1}$$

where:

x is the independent variable

y is the dependent variable

$\bar{x}$ is the mean of the independent variable x

$\bar{y}$ is the mean of the dependent variable y

Based on the covariance, correlation is then calculated by the following formula:

$$r_{(x,y)} = \frac{COV(x,y)}{s_x s_y}$$

where

$r_{(x,y)}$ is the correlation of variables x and y

$COV(x,y)$ is the covariance of variables x and y

s_x is the sample standard deviation of the random variable x

s_y is the sample standard deviation of the random variable y

Correlation standardizes the measure of interdependence between two variables and, consequently, tells you how closely the two variables move. The correlation measurement, called a correlation coefficient, will always take on a value between 1 and -1 :

- *If the correlation coefficient is 1*, the variables have a perfect positive correlation. This means that if one variable moves a given amount, the second moves proportionally in the same direction. A positive correlation coefficient less than one indicates a less than perfect positive correlation, with the strength of the correlation growing as the number approaches one.
- *If correlation coefficient is 0*, no relationship exists between the variables. If one variable moves, you can make no predictions about the movement of the other variable; they are uncorrelated.
- *If correlation coefficient is -1* , the variables are perfectly negatively correlated (or inversely correlated) and move in opposition to each other. If one variable increases, the other variable decreases proportionally. A negative correlation coefficient greater than -1 indicates a less than perfect negative correlation, with the strength of the correlation growing as the number approaches -1 .

The data monitor sampler takes all samples in memory and continually calculates correlation values using this formula. As an example, you could define an event correlation data monitor that displays a correlation between the number of times a network is being reconnoitered, and if that is related to the number of attacks that the network is receiving.

Correlation Rule

A programmed procedure that expresses conditions and actions, and evaluates normal or correlation events. A rule has two parts: a condition and an action.

A condition determines whether a state exists and satisfies related expressions. If so, an action expression defines the response to the condition.

A rule can have one or more conditions. If there is one condition, the rule acts as a filtering tool. If there is more than one condition, the rule acts as a correlation tool. A rule can be created for any incoming event from one or more event generators, with various conditions, logic statements, and thresholds.

See ["Identity Correlation" on page 565](#) for details on using session correlation.

Customers

To support managed security service providers (MSSPs), and larger enterprises that need to track activity related to cost centers or divisions, ArcSight can identify particular customers as the source of specific events.

"Customers" can be any client, tenant, or internal identity scheme you designate.

To use the Customers resource, you first model your existing customer or client base into the Customers resource tree in the Navigator panel, as described in ["Managing Customers" on page 172](#). Then, using the Customer URI values established in the resource tree, you configure the appropriate SmartConnectors to report these values through their Customer URI attributes.

Once you have implemented your Customer resource tree in the ArcSight Console, and configured the relevant SmartConnectors to report these Customer URIs, you can apply the Customer resources as filter condition arguments.

For example, your Customers resource tree might include a branch that translates into a Customer URI of: /All Customers/Brokerages/Central States/Kansas City/Jones&Co.

In the SmartConnectors resource tree, for those connectors that apply to this customer, you would apply this same string as the value for those connectors' Customer URI attribute, found under the Network section of the **Connector: Default: Content** tab. Thereafter, events reported by those connectors can be filtered in or out by referencing that branch of the Customers resource tree.

Dashboards

Dashboards are a graphical display of data gathered from one or more [Data Monitors](#).

Dashboards can display data in a number of graphical formats, including pie and bar charts, tables, and custom layouts.

Data monitors are views within the dashboard that can be configured to report on events, filters, rules, and other data or information that is of particular interest to you. Data monitors can be arranged within dashboards in numerous viewing layouts.

When you right-click in a dashboard, you can choose from options described in ["Dashboard Context Menu Commands" on the next page](#).

Dashboard Context Menu Commands

Option	Description
Save Dashboard	Save any changes you have made to the dashboard and its data monitors.
Save As	Save the configured data monitors and dashboard under a different name.
Close Dashboard	Close the dashboard and remove it from the Viewer panel.
Dashboard>Zoom In / Zoom Out / Fit in	Visually enlarge or reduce the data monitors presented in the Viewer panel. Size the data monitors to allow them to all appear simultaneously in the current Viewer panel.
Data Monitor>Edit	Edit the current data monitor in the Inspect/Edit panel.
Data Monitor>Disable	Turn off the current data monitor.
Data Monitor>Float / Minimize / Close	Float, minimize, or close the current data monitor.
Show>DataMonitorName	Restore minimized data monitors.
Show Details / Show Detailed Channels	Show the event details for the currently selected element in a data monitor graphic, such as a pie chart, or display each of the monitor's elements in detail in separate active channel grids. This is also called "drilling down."
Analyze in Channel	Create an active channel or filter condition based on the highlighted event. The Analyze in Channel command uses the event's attribute type (its column heading), and the particular event's field value (for example, an exact IP address), to formulate filtered channels based on these two factors. The operators can include Create Channel [X = Y] and Add Condition [X = Y] to Editor .
Tools	Choose one of the network tools to explore the origin of the selected event item.
Show Scroll Bar	Toggles a scroll bar on and off in the selected data monitor. Use the scroll bar to show additional rows of tabular data if present.
Export > Data Monitor/Dashboard as ...	Save the selected data monitor or the dashboard in the JPEG (graphic), CSV (comma-separated value or text-based), or HTML file format.

Data Fields

Processed events are composed of several attributes, each of which is a data field with its own characteristics. These event schema data fields fall into the groups shown in the following sections.

Each attribute has both a **Label** that you see in the ArcSight Console and a unique **Script Alias** you use to refer to the attribute in filters, rules, or Velocity templates. The **Data Type** lets you know how to handle the attribute, and the **Default Turbo Level** indicates whether an attribute

is, by default, classified as **1** (essential, or "fastest") or **2** (optional, or "faster"). Turbo Level 3 ("complete") isn't designated because it applies to additional data not represented here.

The easiest way to view all event fields is on the Event Inspector (Event tab) or Common Conditions Editor (CCE) on the Console.

To display the Event Inspector:

1. Select an event in a grid view like an active channel.
2. Right-click and choose **Show event details**.

The event's details appear in the Event Inspector. To view *all* event fields, make sure that no field set is selected to limit the set of fields shown. Select **Clear** from the drop-down menu above the list of event fields. With no field set selected, the drop-down shows "Select a Field Set".



Note: For a list of ArcSight's Common Event Format (CEF) abbreviations, ask your OpenText ArcSight Support representative for the tech note entitled Implementing ArcSight CEF.

Attacker Group

Attacker Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Attacker Group Field Description
Address	attackerAddress	IP address	1	The IP address of the device hosting the attacker. Note: For older versions of ESM, see " Device Custom Group " on page 929 for IPv6 address data type.
Asset ID	attackerAssetId	Resource	2	The asset that represents the device hosting the attacker.
Asset Name	attackerAssetName	String	2	The name of the asset that represents the device hosting the attacker.
Asset Resource	attackerAssetResource	Resource	2	The Resource of the asset that represents the device hosting the attacker.

Attacker Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Attacker Group Field Description
DNS Domain	attackerDnsDomain	String	2	The Domain Name Service domain name associated with the device hosting the attacker.
FQDN	attackerFqdn	String	2	The fully qualified domain name associated with the device hosting the attacker.
Note about Geo fields: Not all IPv6 addresses are mapped to Geo fields described below. In such cases, the corresponding Geo information for the IPv6 address will be blank on the Viewer panel.				
Geo	attackerGeo	GeoDescriptor	1	The geographical information.
Geo Country Code	attackerGeoCountryCode	String	1	The identifier for the national-political state in which a device resides.
Geo Country Flag URL	attackerGeoCountryFlagUrl	String	1	The URL of an image of the flag of the national-political state in which the device resides.
Geo Country Name	attackerGeoCountryName	String	1	The name of the national-political state where a device resides.
Geo Descriptor ID	attackerGeoDescriptorId	ID	1	The internal ID of the geographical reference.
Geo Latitude	attackerGeoLatitude	Double	1	The latitude of a device.
Geo Location Info	attackerGeoLocationInfo	String	2	Other, free-form text information about the device's location.
Geo Longitude	attackerGeoLongitude	Double	1	The Longitude of a device.
Geo Postal Code	attackerGeoPostalCode	String	1	The postal code of the device's location, as assigned by the national-political state where it resides.

Attacker Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Attacker Group Field Description
Geo Region Code	attackerGeoRegionCode	String	1	The identifier of the sub-region of the national-political state where a device resides. The style of the identifier varies with the host country.
Host Name	attackerHostName	String	2	The name of the device hosting the attacker.
MAC Address	attackerMacAddress	MAC address	2	The MAC address associated with the source of the attack (which may or may not be the MAC address of the host device).
NT Domain	attackerNtDomain	String	2	The Windows NT domain associated with the device hosting the attacker.
Port	attackerPort	Integer	1	The network port associated with the source of the attack.
Process ID	attackerProcessId	Integer	2	The ID of the process associated with the source of the attack.
Process Name	attackerProcessName	String	2	The name of process associated with the source of the attack.
Service Name	attackerServiceName	String	2	The name of service associated with the source of the attack.
Translated Address	attackerTranslatedAddress	IP address	1	If network address translation is an issue, this is the translated IP address of the device hosting the attacker.
Translated Port	attackerTranslatedPort	Integer	1	If network address translation is an issue, this is the translated source port associated with the attack. This can happen in a NAT environment.

Attacker Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Attacker Group Field Description
Translated Zone	attackerTranslatedZone	Zone	1	If network address translation is an issue, this is the network zone associated with the translated IP address of the device hosting the attacker.
Translated Zone External ID	attackerTranslatedZoneExternalID	String	1	Returns the external ID for this reference.
Translated Zone ID	attackerTranslatedZoneID	String	1	Returns the ID for the resource in this resource reference.
Translated Zone Name	attackerTranslatedZoneName	String	1	See the common set of resource attributes. It is assumed that the name is always the last field of the URI.
Translated Zone Reference ID	attackerTranslatedZoneReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Translated Zone Resource	attackerTranslatedZoneResource	Resource	1	Locates the resource described by this reference.
Translated Zone URI	attackerTranslatedZoneURI	String	1	Returns the URI for this reference.
User ID	attackerUserId	String	2	The identifier associated with the OS or application of the attacker, at the source of the attack.
User Name	attackerUserName	String	2	The name associated with the attacker, at the source of the attack.
User Privileges	attackerUserPrivileges	String	2	The user-privilege associated with the attacker, at the source of the attack.
Zone	attackerZone	Zone	1	The network zone in which the attacker's device resides.

Attacker Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Attacker Group Field Description
Zone External ID	attackerZoneExternalID	String	1	Returns the external ID for this reference.
Zone ID	attackerZoneID	String	1	Returns the ID for the resource in this resource reference.
Zone Name	attackerZoneName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.
Zone Reference ID	attackerZoneReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Zone Resource	attackerZoneResource	Resource	1	Locates the resource described by this reference.
Zone URI	attackerZoneURI	String	1	See the common set of resource attributes.

Connector Group

This group category falls into the device-to-Manager information chain. The chain begins at **Device**, which is the actual network hardware that senses an event. In cases where data is concentrated or otherwise pre-processed, it may be passed to a trusted reporting Final Device before reaching an **Original Connector**. Although the **Original Connector** is usually the only connector, if the data passes up through a Manager hierarchy, the chain includes handling by **Connector** stages that are the ArcSight Forwarding Connectors that facilitate Manager-to-Manager connections.



Note: Since connectors are not registered to the local Manager, the Original Agent is not known and all the Original Agent fields are therefore blank and do not need to be displayed.

Connector Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Connector Group Field Description
Address	connectorAddress	IP address	1	The IP address of the device hosting the SmartConnector.
Asset ID	connectorAssetId	Resource	1	The asset that represents the device hosting the SmartConnector.
Asset Name	connectorAssetName	String	1	The connector's asset name.
Asset Resource	connectorAssetResource	Resource	1	The connector resource.
Descriptor ID	connectorDescriptorId	ID	1	The connector descriptor.
DNS Domain	connectorDnsDomain	String	1	The Domain Name Service domain name associated with the device hosting the SmartConnector.
Host Name	connectorHostName	String	1	The name of the device hosting the SmartConnector.
ID	connectorId	String	1	The identifier associated with the SmartConnector configuration resource. The format is connectorID(1) connectorID(2) ...
MAC Address	connectorMacAddress	MacAddress	1	The MAC address associated with the SmartConnector (which may or may not be the MAC address of the host device.)
Name	connectorName	String	1	The user-supplied name of the associated SmartConnector configuration resource.
NT Domain	connectorNtDomain	String	1	The Windows NT domain associated with the device hosting the SmartConnector.
Receipt Time	connectorReceiptTime	DateTime	2	The time the event arrived at the SmartConnector.
Severity	connectorSeverity	Connector Severity Enumeration	1	The normalized ArcSight form of the event severity value provided by the SmartConnector.

Connector Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Connector Group Field Description
Time Zone	connectorTimeZone	String	1	The time zone reported by the device hosting the SmartConnector (as TLA).
Time Zone Offset	connectorTimeZoneOffset	Integer	1	The time zone reported by the device hosting the SmartConnector (shown as a UTC offset). Note that device times may be less accurate than other sources.
Translated Address	connectorTranslatedAddress	IP address	1	If network address translation is an issue, this is the translated IP address of the device hosting the SmartConnector.
Translated Zone	connectorTranslatedZone	Zone	1	If network address translation is an issue, this is the Network Zone associated with the translated IP address of the device hosting the SmartConnector.
Translated Zone External ID	connectorTranslatedZoneExternalID	String	1	Returns the external ID for this reference.
Translated Zone ID	connectorTranslatedZoneID	String	1	Returns the ID for the resource in this resource reference.
Translated Zone Name	connectorTranslatedZoneName	String	1	Returns the name from the URI. It assumes that the name is always the last field of the URI.
Translated Zone Reference ID	connectorTranslatedZoneReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference is stored and uniquely identified in the database.
Translated Zone Resource	connectorTranslatedZoneResource	Resource	1	Locates the resource described by this reference.

Connector Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Connector Group Field Description
Translated Zone URI	connectorTranslatedZoneURI	String	1	Returns the URI for this reference.
Type	connectorType	String	1	A description of the type of SmartConnector that reported the event.
Version	connectorVersion	String	1	The software revision number of the SmartConnector that reported the event
Zone	connectorZone	Zone	1	The network zone in which the device hosting this SmartConnector resides.
Zone External ID	connectorZoneExternalID	String	1	Returns the external ID for this reference.
Zone ID	connectorZoneID	String	1	Returns the ID for the resource in this resource reference.
Zone Name	connectorZoneName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.
Zone Reference ID	connectorZoneReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Zone Resource	connectorZoneResource	Resource	1	Locates the resource described by this reference.
Zone URI	connectorZoneURI	String	1	Returns the URI for this reference.

Category Group

Category Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Category Group Field Description
Behavior	categoryBehavior	String	1	Describes the action taken with or by the object.
Custom Format Field	categoryCustomFormatField	String	1	Describes the content of a custom formatted field, if present.
Descriptor ID	categoryDescriptorId	ID	1	The unique ID for the sensor that reported the event
Device Group	categoryDeviceGroup	String	1	The type of event. For example, logging into a firewall is an Operating System type of event.
Device Type	categoryDeviceType	String	2	The type of device. For example, logging into a firewall, would show the Device Type as Firewall.
Object	categoryObject	String	1	Describes the physical or virtual object that was the focus of the event
Outcome	categoryOutcome	String	1	Indicates whether the action was successfully applied to the object.
Significance	categorySignificance	String	1	Characterizes the event from a network-intrusion-detection perspective.
Technique	categoryTechnique	String	1	Describes the method used to apply the action to the object.
Tuple Description	categoryTupleDescription	String	1	The prose description of the event category, assembled from the category components.

Destination Group

Destination Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Destination Group Field Description
Address	destinationAddress	IP address	1	The IP address of the destination device. Note: For older versions of ESM, see " Device Custom Group " on page 929 for IPv6 address data type.
Asset ID	destinationAssetId	Resource	2	The asset that represents the device that was the network traffic's destination.
Asset Name	destinationAssetName	String	2	The name of the device.
Asset Resource	destinationAssetResource	Resource	2	See the common set of resource attributes.
DNS Domain	destinationDnsDomain	String	2	The Domain Name Service domain name associated with the user at the destination device.
FQDN	destinationFqdn	String	2	The fully qualified domain name associated with the destination device.
Note about Geo fields: Not all IPv6 addresses are mapped to Geo fields described below. In such cases, the corresponding Geo information for the IPv6 address will be blank on the Viewer panel.				
Geo	destinationGeo	GeoDescriptor	1	See the common set of geographical attributes.
Geo Country Code	destinationGeoCountryCode	String	1	The identifier for the national-political state in which a device resides.
Geo Country Flag URL	destinationGeoCountryFlagUrl	String	1	The URL of an image of the flag of the national-political state in which the device resides.

Destination Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Destination Group Field Description
Geo Country Name	destinationGeoCountryName	String	1	The name of the national-political state where a device resides.
Geo Descriptor ID	destinationGeoDescriptorId	ID	1	The internal ID of the geographical reference.
Geo Latitude	destinationGeoLatitude	Double	1	The destination latitude of the device.
Geo Location Info	destinationGeoLocationInfo	String	1	Other, free-form text information about the device's location.
Geo Longitude	destinationGeoLongitude	Double	1	The destination longitude.
Geo Postal Code	destinationGeoPostalCode	String	1	The postal code of the device's location, as assigned by the national-political state where it resides.
Geo Region Code	destinationGeoRegionCode	String	1	The identifier of the sub-region of the national-political state where a device resides. The style of the identifier varies with the host country.
Host Name	destinationHostName	String	2	The name of the destination device.
MAC Address	destinationMacAddress	MAC address	2	The MAC address associated with the network traffic's destination (which may or may not be the MAC address of the host device).
NT Domain	destinationNtDomain	String	2	The Windows NT domain associated with the destination device.
Port	destinationPort	Integer	1	The network port associated with the network traffic's destination.

Destination Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Destination Group Field Description
Process ID	destinationProcessId	Integer	2	The ID of the process associated with the network traffic's destination.
Process Name	destinationProcessName	String	2	The name of the process associated with the network traffic's destination.
Service Name	destinationServiceName	String	2	The name of service associated with the network traffic's destination.
Translated Address	destinationTranslatedAddress	IP address	1	If network address translation is an issue, this is the translated IP address of the device that was the network traffic's destination.
Translated Port	destinationTranslatedPort	Integer	1	If network address translation is an issue, this is the translated source port associated with the attack.
Translated Zone	destinationTranslatedZone	Zone	1	If network address translation is an issue, this is the network zone associated with the translated IP address of the device at the network's traffic's destination.
Translated Zone External ID	destinationTranslatedZoneExternalID	String	1	Returns the external ID for this reference.
Translated Zone ID	destinationTranslatedZoneID	String	1	Returns the ID for the resource in this resource reference.
Translated Zone Name	destinationTranslatedZoneName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.

Destination Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Destination Group Field Description
Translated Zone Reference	destinationTranslatedZoneReferenceID	ID	1	See the common set of resource attributes.
Translated Zone Resource	destinationTranslatedZoneResource	Resource	1	Locates the resource described by this reference.
Translated Zone URI	destinationTranslatedZoneURI	String	1	Returns the URI for this reference.
User ID	destinationUserId	String	2	The OS- or application-based identifier associated with the user at the network traffic's destination.
User Name	destinationUserName	String	2	The name associated with the user at the network traffic's destination.
User Privileges	destinationUserPrivileges	String	2	The privileges accorded the user at the network traffic destination.
Zone	destinationZone	Zone	1	The network zone in which the destination device resides.
Zone External ID	destinationZoneExternalID	String	1	Returns the external ID for this reference.
Zone ID	destinationZoneID	String	1	Returns the ID for the resource in this resource reference.
Zone Name	destinationZoneName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.

Destination Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Destination Group Field Description
Zone Reference ID	destinationZoneReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Zone Resource	destinationZoneResource	Resource	1	Locates the resource described by this reference.
Zone URI	destinationZoneURI	String	1	See the common set of resource attributes.

Device Group

This category falls into the device-to-Manager information chain. The chain begins at **Device**, which is the actual network hardware that senses an event. In cases where data is concentrated or otherwise pre-processed, it may be passed to a trusted reporting **Final Device** before reaching an **Original Connector**. Although the **Original Connector** is usually the only connector, if the data passes up through a Manager hierarchy the chain includes handling by **Connector** stages that are the Manager SmartConnectors that facilitate Manager-to-Manager connections.

Device Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Device Group Field Description
Action	deviceAction	String	2	The device-specific description of some activity associated with the event
Address	deviceAddress	IP address	1	The IP address of the device hosting the sensor. Note: For older versions of ESM, see " Device Custom Group " on page 929 for IPv6 address data type.

Device Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Device Group Field Description
Asset ID	deviceAssetId	Resource	1	The asset that represents the device hosting the sensor.
Asset Name	deviceAssetName	String	1	The name of the device.
Asset Resource	deviceAssetResource	Resource	1	The resource the asset represents.
Descriptor ID	deviceDescriptorId	ID	1	The asset's descriptor ID.
Direction	deviceDirection	Device Direction Enumeration	2	Whether the traffic was inbound or outbound.
DNS Domain	deviceDnsDomain	String	1	The Domain Name Service domain name associated with the device hosting the sensor.
Domain	deviceDomain	String	2	The specific domain containing the sensor device associated with the event
Event Category	deviceEventCategory	String	2	The category description included with the event as reported by the device.
Event Class ID	deviceEventClassId	String	2	The device-specific identifier associated with this type of event Note: A generic UNIX syslog parser displays the ID in this format: arcsight:x:x
External ID	deviceExternalId	String	1	The external identifier associated with this sensor device, if provided by the vendor.
Facility	deviceFacility	String	1	The sensor submodule that reported the event
Host Name	deviceHostName	String	1	The name of the device hosting the sensor.

Device Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Device Group Field Description
Inbound Interface	deviceInboundInterface	String	1	The NIC card on the sensor device that received the network traffic associated with the event.
MAC Address	deviceMacAddress	MAC address	1	The MAC address associated with the source of the attack (which may or may not be the MAC address of the host device).
NT Domain	deviceNtDomain	String	1	The Windows NT domain associated with the device hosting the sensor.
Outbound Interface	deviceOutboundInterface	String	1	The NIC card on the sensor device that transmitted the network traffic associated with the event.
Payload ID	devicePayloadId	String	2	The internal identifier associated with a payload object associated with this event.
Process ID	deviceProcessId	Integer	2	The ID of the sensor device process that reported the event.
Process Name	deviceProcessName	String	1	The name of the sensor device process that reported the event.
Product	deviceProduct	String	1	The product name of the sensor device.
Receipt Time	deviceReceiptTime	DateTime	2	The time when the sensor device observed the event.

Device Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Device Group Field Description
Severity	deviceSeverity	String	2	The device-specific assessment of event severity. This assessment varies with the device involved.
Time Zone	deviceTimeZone	String	1	The time zone reported by the device hosting the sensor device (shown as TLA).
Time Zone Offset	deviceTimeZoneOffset	Integer	1	The time zone reported by the device hosting this sensor device (shown as an offset from UTC).
Translated Address	deviceTranslatedAddress	IP address	1	If network address translation is an issue, this is the translated IP address of the device hosting the sensor.
Translated Zone	deviceTranslatedZone	Zone	1	If network address translation is an issue, this is the network zone associated with the translated IP address of the device hosting the sensor.
Translated Zone External ID	deviceTranslatedZoneExternalID	String	1	Returns the external ID for this reference.
Translated Zone ID	deviceTranslatedZoneID	String	1	Returns the ID for the resource in this resource reference.
Translated Zone Name	deviceTranslatedZoneName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.

Device Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Device Group Field Description
Translated Zone Reference ID	deviceTranslatedZoneReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Translated Zone Resource	deviceTranslatedZoneResource	Resource	1	Locates the resource described by this reference.
Translated Zone URI	deviceTranslatedZoneURI	String	1	Returns the URI for this reference.
Vendor	deviceVendor	String	1	The vendor who manufactured or sold the sensor device.
Version	deviceVersion	String	1	The software revision number of the sensor device.
Zone	deviceZone	Zone	1	The network zone in which the sensor's device resides.
Zone External ID	deviceZoneExternalID	String	1	Returns the external ID for this reference.
Zone ID	deviceZoneID	String	1	Returns the ID for the resource in this resource reference.
Zone Name	deviceZoneName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.

Device Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Device Group Field Description
Zone Reference ID	deviceZoneReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been persisted and given a unique database identifier.
Zone Resource	deviceZoneResource	Resource	1	Locates the resource described by this reference.
Zone URI	deviceZoneURI	String	1	See the common set of resource attributes.

Device Custom Group

Device Custom Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Device Custom Group Field Description
Date1	deviceCustomDate1	DateTime	2	First customDate
Date1 Label	deviceCustomDate1Label	String	2	First customDate label
Date2	deviceCustomDate2	DateTime	2	Second customDate
Date2 Label	deviceCustomDate2Label	String	2	Second customDate label
Number1	deviceCustomNumber1	Long	2	First customNumber
Number1 Label	deviceCustomNumber1Label	String	2	First customNumber label
Number2	deviceCustomNumber2	Long	2	Second customNumber
Number2 Label	deviceCustomNumber2Label	String	2	Second customNumber label

Device Custom Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Device Custom Group Field Description
Number3	deviceCustomNumber3	Long	2	Third customNumber
Number3 Label	deviceCustomNumber3Label	String	2	Third customNumber label
String1	deviceCustomString1	String	2	First customString
String1 Label	deviceCustomString1Label	String	2	First customString label
String2	deviceCustomString2	String	2	Second customString
String2 Label	deviceCustomString2Label	String	2	Second customString label
String3	deviceCustomString3	String	2	Third customString
String3 Label	deviceCustomString3Label	String	2	Third customString label
String4	deviceCustomString4	String	2	Fourth customString
String4 Label	deviceCustomString4Label	String	2	Fourth customString label
String5	deviceCustomString5	String	2	Fifth customString
String5 Label	deviceCustomString5Label	String	2	Fifth customString label
String6	deviceCustomString6	String	2	Sixth customString
String6 Label	deviceCustomString6Label	String	2	Sixth customString label
Floating Point1	deviceCustomFloatingPoint1	String	2	First custom floating point

Device Custom Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Device Custom Group Field Description
Floating Point1 Label	deviceCustomFloatingPoint1Label	Double	2	First custom floating point label
Floating Point2	deviceCustomFloatingPoint2	String	2	Second custom floating point
Floating Point2 Label	deviceCustomFloatingPoint2Label	Double	2	Second custom floating point label
Floating Point3	deviceCustomFloatingPoint3	String	2	Third custom floating point
Floating Point3	deviceCustomFloatingPoint3Label	Double	2	Third custom floating point label
Floating Point4	deviceCustomFloatingPoint4	String	2	Fourth custom floating point
Floating Point4 Label	deviceCustomFloatingPoint4Label	Double	2	Fourth custom floating point label
Note: For versions prior to ESM 6.11.0, the following device custom fields are still available to support IPv6 addresses. The IPv6 address is stored in the database in its full form, and ESM displays it in a simplified format, if applicable. If you used the following fields for IPv4 addresses, they are displayed as IPv6-embedded IPv4 addresses. You should also check your SmartConnector version and accompanying documentation. Old connectors will continue to use these fields for IPv6 addresses. Updated connectors will send IPv6 addresses to the appropriate address fields.				
IPv6 Address1	deviceCustomIPv6Address1	IPv6 address	2	First custom IPV6 address
IPv6 Address1 Label	deviceCustomIPv6Address1Label	String	2	First custom IPV6 address label
IPv6 Address2	deviceCustomIPv6Address2	IPv6 address	2	Second custom IPV6 address
IPv6 Address2 Label	deviceCustomIPv6Address2Label	String	2	Second custom IPV6 address label
IPv6 Address3	deviceCustomIPv6Address3	IPv6 address	2	Third custom IPV6 address

Device Custom Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Device Custom Group Field Description
IPv6 Address3 Label	deviceCustomIPv6Address3Label	String	2	Third custom IPv6 address label
IPv6 Address4	deviceCustomIPv6Address4	IPv6 address	2	Fourth custom IPv6 address
IPv6 Address4 Label	deviceCustomIPv6Address4Label	String	2	Fourth custom IPv6 address label

Event Group

Event Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Event Group Field Description
Additional Data	additionalData	AdditionalData	3	Reference to additional data.
Aggregated Event Count	(not applicable)	(not applicable)	N/A	A derived field that reports the number of actual events collectively represented by the event in question.
Application Protocol	applicationProtocol	String	2	A description of the application layer protocol. May be set, but defaults to Target Port lookup (FTP).
Base Event IDs	baseEventIds	ID	2	The array of event IDs that contributed to generating this correlation event. This is populated only in correlated events.
Bytes In	bytesIn	Integer	2	Number of bytes transferred into the device during this transaction (this would typically be associated with entries in HTTP logs).

Event Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Event Group Field Description
Bytes Out	bytesOut	Integer	2	Number of bytes transferred out of the device during this transaction (this would typically be associated with entries in HTTP logs).
Concentrator Connectors	concentratorConnectors	ConnectorDescriptor	2	The chain of concentrators that forwarded the event. This is not yet exposed in the user interface.
Concentrator Devices	concentratorDevices	DeviceDescriptor	2	The list of devices that concentrate events, if applicable. This is not exposed in the user interface.
Correlated Event Count	(not applicable)	(not applicable)	N/A	A derived field that reports the number of actual events that had to occur to cause a correlation event to occur.
Crypto Signature	cryptoSignature	String	2	The signature of the event object (meaning in this alert, as opposed to the occurrence represented by the event). Not yet supported.
Customer	customer	Customer	1	The "customer" resource reference. This is used in MSSP environments to describe the client or divisional entity to whom the event applies.
Customer External ID	customerExternalID	String	1	Returns the external ID for this reference.
Customer ID	customerID	String	1	Returns the ID for the resource in this resource reference.

Event Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Event Group Field Description
Customer Name	customerName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.
Customer Reference ID	customerReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Customer Resource	customerResource	Resource	1	Locates the resource described by this reference.
Customer URI	customerURI	String	1	Returns the URI for this reference.
End Time	endTime	DateTime	1	Event ends (defaults to deviceReceiptTime).
Event ID	eventId	ID	1	Long 64-bit value identifying an event. A negative event ID is normal. The less significant 48 bits are assigned to a newly received event by the receiving Manager; these bits uniquely identify the event in the database of that Manager. The more-significant 16 bits are used to store forwarding information. When an event ID with 1 in the topmost bit is represented as Java long value, the event ID value is interpreted as a negative number according to JVM rules. When displayed, such an event ID appears as a decimal number with a - (minus) sign in front of it.

Event Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Event Group Field Description
Event Outcome	eventOutcome	String	2	The outcome of the event as reported by the device (when applicable). For example, Windows reports an event as audit_success or audit_failure.
External ID	externalId	String	2	A reference to the ID used by an external device. This is useful for tracking devices that create events that contain references to these IDs (for example, ManHunt).
Generator	generator	null	1	The "generator" resource reference (the resource that generated the event. This is the subcomponent in the connector that generates the event.
Generator External ID	generatorExternalID	String	1	Returns the external ID for this reference.
Generator ID	generatorID	String	1	Returns the ID for the resource in this resource reference.
Generator Name	generatorName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.
Generator Reference ID	generatorReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Generator Resource	generatorResource	Resource	1	Locates the resource described by this reference.
Generator URI	generatorURI	String	1	Returns the URI for this reference.

Event Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Event Group Field Description
Locality	locality	LocalityEnumeration	2	The locality associated with the event. Possible values: Local = 0: Events were sent to a Manager from SmartConnectors Forwarded = 1: Events were sent to a Manager from a Forwarding Connector. Remote = 2: Events were fetched from a remote Manager (to display a rule chain, for example) ESMPassThrough = 3: Event is processed by ESM. DirectPassToLogger = 4: Event is not processed by ESM and is passed to Logger or any other product.
Message	message	String	2	A brief comment associated with this event.
Name	name	String	1	An arbitrary string that describes this type of event. Event details included in other parts of an event shouldn't be used in the event name.
Originator	originator	OriginatorEnumeration	1	Holds the value of Source Destination. This determines whether source and destination should be translated to attacker and target or they should be reversed.

Event Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Event Group Field Description
Persistence	persistence	PersistenceEnumeration	2	There are two states: Persisted or Transient. Events default to being Transient and are marked as Persisted as soon as they reach the Batch Alert Persister or when they are loaded by the Alert Broker.
Raw Event	rawEvent	String	1	The original log entry reported by the sensor (synthesized when the sensor does not log to a file or text stream).
Reason	reason	String	2	The cause of the event when applicable. For example, Invalid Password
Rule Thread ID	ruleThreadId	String	2	A single rule can issue many events, based on several triggers, starting with On First Event and ending with On Threshold Timeout. All such events for a single Rule and a single Group By tuple is marked with the same identifier using this attribute.
Session ID	sessionId	Long	2	Tags for events created by a correlation simulation, as part of a particular simulation.
Start Time	startTime	DateTime	1	Event begins (defaults to deviceReceiptTime).
Transport Protocol	transportProtocol	String	1	The format of the transmitted data associated with the event from a network transport perspective (for example, TCP, UDP).
Type	type	TypeEnumeration	1	One of the event types: Base, Correlation, Aggregated, or Action.

Event Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Event Group Field Description
Vulnerability	vulnerability	Vulnerability	2	The vulnerability resource that represents the vulnerability or exposure that may be exploited by this event and is present on the targeted device according to our network model.
Vulnerability External ID	vulnerabilityExternalID	String	2	Returns the external ID for this reference.
Vulnerability ID	vulnerabilityID	String	2	Returns the ID for the resource in this resource reference.
Vulnerability Name	vulnerabilityName	String	2	Returns the name from the URI, which is always assumed to be the last field of the URI.
Vulnerability Reference ID	vulnerabilityReferenceID	ID	2	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Vulnerability Resource	vulnerabilityResource	Resource	2	Locates the resource described by this reference.
Vulnerability URI	vulnerabilityURI	String	2	Returns the URI for this reference.

Event Annotation Group

Event Annotation Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Event Annotation Group Field Description
Audit Trail	eventAnnotationAuditTrail	String	2	The text log of annotation changes. Changes are recorded as sets of comma-separated-value entries.
Comment	eventAnnotationComment	String	2	A text description of the event or associated information.
End Time	eventAnnotationEndTime	DateTime	2	The timestamp for an event annotation.
Event ID	eventAnnotationEventId	ID	2	The event ID for the annotation event.
Flags	eventAnnotationFlags	FlagsValue Set	2	The state of the collaboration flags. Note: The following event fields: isReviewed, Closed, Hidden, Correlated, inCase, hasAction, and Forwarded, are derived from eventAnnotationFlags. You cannot add those individual event fields into a field set, but you can add the eventAnnotationFlags fields instead, then use a local or global variable to specify the desired field. See "Field Set Editor: Local Variables Tab" on page 605 and "Global Variables" on page 609.
Manager Receipt Time	eventAnnotationManagerReceiptTime	DateTime	2	The time the Manager received the event annotation.
Modification Time	eventAnnotationModificationTime	DateTime	2	The time the annotation was modified.

Event Annotation Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Event Annotation Group Field Description
Modified By	eventAnnotationModifiedBy	User	2	The user ID of the person who last edited this annotation.
Modified By External ID	eventAnnotationModifiedByExternalID	String	2	Returns the external ID for this reference.
Modified By ID	eventAnnotationModifiedByID	String	2	Returns the ID for the resource in this resource reference.
Modified By Name	eventAnnotationModifiedByName	String	2	Returns the name from the URI (the last field of the URI).
Modified By Reference ID	eventAnnotationModifiedByReferenceID	ID	2	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Modified By Resource	eventAnnotationModifiedByResource	Resource	2	Locates the resource described by this reference.
Modified By URI	eventAnnotationModifiedByURI	String	2	Returns the URI for this reference.
Stage	eventAnnotationStage	Stage	2	The current disposition of the event. This enables annotation workflow.
Stage Event ID	eventAnnotationStageEventId	ID	2	The reference to an internal identifier for another event. It is used by 'Mark Similar'.
Stage External ID	eventAnnotationStageExternalID	String	2	Returns the external ID for this reference.
Stage ID	eventAnnotationStageID	String	2	Returns the ID for the resource in this resource reference.
Stage Name	eventAnnotationStageName	String	2	Returns the name from the URI, which is always assumed to be the last field of the URI.

Event Annotation Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Event Annotation Group Field Description
Stage Reference ID	eventAnnotationStageReferenceID	ID	2	Returns the unique descriptor ID for this reference. This is populated only if this reference is stored and uniquely identified in the database.
Stage Resource	eventAnnotationStageResource	Resource	2	Locates the resource described by this reference.
Stage Update Time	eventAnnotationStageUpdateTime	ID	2	The time of the last stage change (in UTC).
Stage URI	eventAnnotationStageURI	String	2	Returns the URI for this reference.
Stage User	eventAnnotationStageUser	User	2	The user associated with the current stage. This implements assignment within workflow.
Stage User External ID	eventAnnotationStageUserExternalID	String	2	Returns the external ID for this reference.
Stage User ID	eventAnnotationStageUserID	String	2	Returns the ID for the resource in this resource reference.
Stage User Name	eventAnnotationStageUserName	String	2	Returns the name from the URI, which is always assumed to be the last field of the URI.
Stage User Reference ID	eventAnnotationStageUserReferenceID	ID	2	Returns the unique descriptor ID for this reference. This is populated only if this reference is stored and uniquely identified in the database.
Stage User Resource	eventAnnotationStageUserResource	Resource	2	Locates the resource described by this reference.
Stage User URI	eventAnnotationStageUserURI	String	2	Returns the URI for this reference.
Version	eventAnnotationVersion	Integer	2	The editing version number which increments with each change. This enables optimistic locking.

File Group

File Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	File Group Field Description
Create Time	fileCreateTime	DateTime	2	The time the file was created (in UTC).
Hash	fileHash	String	2	The hash code associated with the file's contents (for example, MD5).
ID	fileId	String	2	The external identifier associated with the file.
Modification Time	fileModificationTime	DateTime	2	The time the file was last changed (in UTC).
Name	fileName	String	2	The name of the file.
Path	filePath	String	2	The directory path to the file in the file system.
Permission	filePermission	String	2	The user permissions associated with the file (sensor specific).
Size	fileSize	Long	2	The size of the file's contents (typically in bytes; sensor specific).
Type	fileType	String	2	The type of file contents (sensor specific).

Final Device Group

This category falls into the device-to-Manager information chain. The chain begins at **Device**, which is the actual network hardware that senses an event. In cases where data is concentrated or otherwise pre-processed, it may be passed to a trusted reporting **Final Device** before reaching an **Original Connector**. Although the **Original Connector** is usually the only connector, if the data passes up through a Manager hierarchy the chain includes handling by **Connector** stages that are the Manager SmartConnectors that facilitate Manager-to-Manager connections.

Final Device Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Final Device Group Field Description
Address	finalDeviceAddress	IP address	2	The IP address of the trusted reporting device. Note: For older versions of ESM, see "Device Custom Group" on page 929 for IPv6 address data type.
Asset ID	finalDeviceAssetId	Resource	2	The asset that represents the trusted reporting device.
Asset Name	finalDeviceAssetName	String	2	The name of the trusted reporting device.
Asset Resource	finalDeviceAssetResource	Resource	2	The resource represented by the trusted reporting device.
Descriptor ID	finalDeviceDescriptorId	ID	2	The descriptor ID of the trusted reporting device.
DNS Domain	finalDeviceDnsDomain	String	2	The Domain Name Service domain name associated with the trusted reporting device.
External ID	finalDeviceExternalId	String	2	The external ID for the trusted reporting device, if provided by the vendor.
Facility	finalDeviceFacility	String	2	A facility or capability of a device. This accommodates concentrators (for example, like syslog, which has a concept of device logging for "parts" of a device).
Host Name	finalDeviceHostName	String	2	The host name of the trusted reporting device.
Inbound Interface	finalDeviceInboundInterface	String	2	The NIC card on the sensor device that received the network traffic associated with the event.
MAC address	finalDeviceMacAddress	MAC address	2	The MAC address associated with the trusted reporting device.
NT Domain	finalDeviceNtDomain	String	2	The Windows NT domain associated with the trusted reporting device.
Outbound Interface	finalDeviceOutboundInterface	String	2	The NIC card on the trusted reporting device.
Process Name	finalDeviceProcessName	String	2	The process name of the trusted reporting device.
Product	finalDeviceProduct	String	2	The product name of the trusted reporting device.

Final Device Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Final Device Group Field Description
Time Zone	finalDeviceTimeZone	String	2	The time zone reported by the trusted reporting device.
Time Zone Offset	finalDeviceTimeZoneOffset	Integer	2	Returns the raw time-zone offset for the trusted reporting device. Note that connector and device times are not always reliably accurate.
Translated Address	finalDeviceTranslated Address	IP address	2	If network address translation is an issue, this is the translated IP address of the trusted reporting device.
Translated Zone	finalDeviceTranslatedZone	Zone	2	If network address translation is an issue, this is the network zone associated with the translated IP address of the trusted reporting device.
Translated Zone External ID	finalDeviceTranslatedZone ExternalID	String	2	Returns the external ID for this reference.
Translated Zone ID	finalDeviceTranslatedZoneID	String	2	Returns the ID for the resource in this resource reference.
Translated Zone Name	finalDeviceTranslatedZone Name	String	2	Returns the name from the URI, which is always assumed to be the last field of the URI.
Translated Zone Reference ID	finalDeviceTranslatedZone ReferenceID	ID	2	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Translated Zone Resource	finalDeviceTranslatedZone Resource	Resource	2	Locates the resource described by this reference.
Translated Zone URI	finalDeviceTranslatedZone URI	String	2	Returns the URI for this reference.
Vendor	finalDeviceVendor	String	2	Device vendor.
Version	finalDeviceVersion	String	2	The software revision number of the trusted reporting device.
Zone	finalDeviceZone	Zone	2	The network zone in which the trusted reporting device resides.
Zone External ID	finalDeviceZoneExternalID	String	2	Returns the external ID for this reference.

Final Device Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Final Device Group Field Description
Zone ID	finalDeviceZoneID	String	2	Returns the ID for the resource in this resource reference.
Zone Name	finalDeviceZoneName	String	2	Returns the name from the URI, which is always assumed to be the last field of the URI.
Zone Reference ID	finalDeviceZoneReferenceID	ID	2	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Zone Resource	finalDeviceZoneResource	Resource	2	Locates the resource described by this reference.
Zone URI	finalDeviceZoneURI	String	2	Returns the URI for this reference.

Flex Group

Flex Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Flex Group Field Description
Date1	flexDate1	DateTime	2	First flex Date.
Date1 Label	flexDate1Label	String	2	Label of first flex Date.
Number1	flexNumber1	Long	2	First flex Number.
Number1 Label	flexNumber1Label	String	2	Label of the first Flex Number.
Number2	flexNumber2	Long	2	Second flex Number.
Number2 Label	flexNumber2Label	String	2	Label of the second Flex Number.
String1	flexString1	String	2	First flex String
String1 Label	flexString1Label	String	2	Label of the first Flex String.
String2	flexString2	String	2	Second flex String.
String2 Label	flexString2Label	String	2	Label of the second Flex String.

Geographical Attributes

Not all IPv6 addresses are mapped to fields described below. In such cases, the corresponding Geo information for the IPv6 address will be blank on the Viewer panel.

Geographical Attributes Data Fields

Attribute Suffix	Description
Descriptor ID	The internal ID of the geographical reference.
Country Code	The identifier for the national-political state in which a device resides.
Country Flag URL	The URL of an image of the flag of the national-political state in which the device resides.
Country Name	The name of the national-political state where a device resides.
Latitude	The latitude of a device.
Location Info	Other, free-form text information about the device's location.
Longitude	The longitude of a device.
Postal Code	The postal code of the device's location, as assigned by the national-political state where it resides.
Region Code	The identifier of the sub-region of the national-political state where a device resides. The style of the identifier varies with the host country.

Manager Group

Label	Script Alias	Data Type	Default Turbo Level	Manager Group Field Description
Receipt Time	managerReceiptTime	DateTime	1	The time at which the Manager first received the event.

Old File Group

Old File Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Old File Group Field Description
Create Time	oldFileCreateTime	DateTime	2	The time the file was created (in UTC).
Hash	oldFileHash	String	2	The hashcode associated with the file's contents (for example, MD5).
ID	oldFileId	String	2	The external identifier associated with the file.
Modification Time	oldFileModificationTime	DateTime	2	The time the file was last changed (in UTC).

Old File Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Old File Group Field Description
Name	oldFileName	String	2	The file's name.
Path	oldFilePath	String	2	The directory path to the file in the file system.
Permission	oldFilePermission	String	2	The user permissions associated with the file (sensor specific).
Size	oldFileSize	Long	2	The size of the file's contents (typically in bytes; sensor specific).
Type	oldFileType	String	2	The type of the file's contents (sensor specific).

Original Connector Group

This category falls into the device-to-Manager information chain. The chain begins at **Device**, the actual network hardware that senses an event. Where data is concentrated or otherwise pre-processed, it may be passed to a trusted reporting **Final Device** before reaching an **Original Connector**. Although the **Original Connector** is usually the only connector, if the data passes up through a Manager hierarchy, the chain includes handling by **Connector** stages that are the Forwarding Connectors that facilitate Manager-to-Manager connections.

Original Connector Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Original Connector Group Field Description
Address	originalConnectorAddress	IP address	2	The IP address of the device hosting the first reporting SmartConnector. Note: For older versions of ESM, see "Device Custom Group" on page 929 for IPv6 address data type.
Asset ID	originalConnectorAssetID	Resource	2	The asset that represents the device hosting the first reporting SmartConnector.
Asset Name	originalConnectorAsset Name	String	2	The first reporting connector's asset name.

Original Connector Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Original Connector Group Field Description
Asset Resource	originalConnectorAsset Resource	Resource	2	The first reporting connector's resource.
Descriptor ID	originalConnectorDescriptorId	ID	2	The first reporting connector's descriptor.
DNS Domain	originalConnectorDns Domain	String	2	The Domain Name Service domain name associated with the device hosting the first reporting SmartConnector.
Host Name	originalConnectorHostName	String	2	The name of the device hosting the first reporting SmartConnector.
ID	originalConnectorId	String	2	The ID of the connector. The format is connectorId(1) connectorId(2) ...
MAC address	originalconnectorMac Address	MAC address	2	The MAC address associated with the first reporting SmartConnector (which may or may not be the MAC address of the host device.)
Name	originalconnectorName	String	2	User-supplied name of the first reporting connector.
NT Domain	originalconnectorNtDomain	String	2	The Windows NT domain associated with the device hosting the first reporting SmartConnector.
Time Zone	originalconnectorTimeZone	String	2	The time zone reported by the device hosting the first reporting SmartConnector.
Time Zone Offset	originalconnectorTimeZoneOffset	Integer	2	Returns the raw time-zone offset for the first reporting connector's time zone. Note that device and connector times may not be reliably accurate.

Original Connector Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Original Connector Group Field Description
Translated Address	originalconnectorTranslatedAddress	IP address	2	If network address translation is an issue, this is the translated IP address of the device hosting the first reporting SmartConnector.
Translated Zone	originalconnectorTranslatedZone	Zone	2	If network address translation is an issue, this is the Network Zone associated with the translated IP address of the device hosting the first reporting SmartConnector.
Translated Zone External ID	originalconnectorTranslatedZoneExternalID	String	2	Returns the external ID for this reference.
Translated Zone ID	originalconnectorTranslatedZoneID	String	2	Returns the ID for the resource in this resource reference.
Translated Zone Name	originalconnectorTranslatedZoneName	String	2	Returns the name from the URI, which is always assumed to be the last field of the URI.
Translated Zone Reference ID	originalconnectorTranslatedZoneReferenceID	ID	2	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Translated Zone Resource	originalconnectorTranslatedZoneResource	Resource	2	Locates the resource described by this reference.
Translated Zone URI	originalconnectorTranslatedZoneURI	String	2	Returns the URI for this reference.
Type	originalconnectorType	String	2	A string that describes the type of the first reporting connector. This is not the same as the device type.

Original Connector Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Original Connector Group Field Description
Version	originalconnectorVersion	String	2	The software revision number of the SmartConnector that first reported the event.
Zone	originalconnectorZone	Zone	2	The network zone in which the device hosting the first reporting SmartConnector resides.
Zone External ID	originalconnectorZone ExternalID	String	2	Returns the external ID for this reference.
Zone ID	originalconnectorZoneID	String	2	Returns the ID for the resource in this resource reference.
Zone Name	originalconnectorZoneName	String	2	Returns the name from the URI, which is always assumed to be the last field of the URI.
Zone Reference ID	originalconnectorZone ReferenceID	ID	2	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and is uniquely identified in the database.
Zone Resource	originalconnectorZone Resource	Resource	2	Locates the resource described by this reference.
Zone URI	originalconnectorZoneURI	String	2	Returns the URI for this reference.

Request Group

Request Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Request Group Field Description
Client Application	requestClientApplication	String	2	The client application (such as a web browser) used to issue the request.
Client Application	requestClientApplication	String	2	A description of the client application used to initiate this request, for example, the HTTP User connector.
Context	requestContext	String	2	A description of the content from which the request originated, for example, the HTTP Referrer.
Cookies	requestCookies	String	2	Cookie data offered by the client application as part of the request.
Method	requestMethod	String	2	The style of the request, that is, for an HTTP request this could be PUT or GET.
Protocol	requestProtocol	String	2	The communication protocol used when issuing the request.
URL	requestUrl	String	2	A universal resource locator associated with the event.
URL Authority	requestUrlAuthority	String	2	The URL component used for authentication and authorization.
URL File Name	requestUrlFileName	String	2	The URL component that refers to the file containing the resource.
URL Host	requestUrlHost	String	2	The URL component that specifies the host device where the resource resides.
URL Port	requestUrlPort	Integer	2	The URL component that specifies the port to contact on the host device where the resource resides.
URL Query	requestUrlQuery	String	2	The URL component that specifies the query to use to request the resource.

Source Group

Source Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Source Group Field Description
Address	sourceAddress	IP address	1	The IP address of the source device. Note: For older versions of ESM, see "Device Custom Group" on page 929 for IPv6 address data type.
Asset ID	sourceAssetId	Resource	2	The asset that represents the device that was the network traffic's source.
Asset Name	sourceAssetName	String	2	The name of the device.
Asset Resource	sourceAssetResource	Resource	2	See the common set of resource attributes.
DNS Domain	sourceDnsDomain	String	2	The Domain Name Service domain name associated with the user at the source device.
FQDN	sourceFqdn	String	2	The fully qualified domain name associated with the source device. This has no value if either the host name or DNS domain are without a value.
Geo	sourceGeo	GeoDescriptor	1	The geographical information.
Note about Geo fields: Not all IPv6 addresses are mapped to Geo fields described below. In such cases, the corresponding Geo information for the IPv6 address will be blank on the Viewer panel.				
Geo Country Code	sourceGeoCountryCode	String	1	The identifier for the national-political state in which a device resides.
Geo Country Flag URL	sourceGeoCountryFlagUrl	String	1	The URL of an image of the flag of the national-political state in which the device resides.
Geo Country Name	sourceGeoCountryName	String	1	The name of the national-political state where a device resides.
Geo Descriptor ID	sourceGeoDescriptorId	ID	1	The internal ID of the geographical reference.

Source Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Source Group Field Description
Geo Latitude	sourceGeoLatitude	Double	1	The latitude of a device.
Geo Location Info	sourceGeoLocationInfo	String	1	Other, free-form text information about the device's location.
Geo Longitude	sourceGeoLongitude	Double	1	The Longitude of a device.
Geo Postal Code	sourceGeoPostalCode	String	1	The postal code of the device's location, as assigned by the national-political state where it resides.
Geo Region Code	sourceGeoRegionCode	String	1	The identifier of the sub-region of the national-political state where a device resides. The style of the identifier varies with the host country.
Host Name	sourceHostName	String	2	The name of the source device.
MAC Address	sourceMacAddress	MAC address	2	The MAC address associated with the network traffic's source (which may or may not be the MAC address of the host device).
NT Domain	sourceNtDomain	String	2	The Windows NT domain associated with the source device.
Port	sourcePort	Integer	1	The network port associated with the network traffic's source.
Process ID	sourceProcessId	Integer	2	The ID of the process associated with the source of the network traffic.
Process Name	sourceProcessName	String	2	The name of the process associated with the source of the network traffic.
Service Name	sourceServiceName	String	2	The name of the service associated with the network traffic's source.
Translated Address	sourceTranslatedAddress	IP address	1	If network address translation is an issue, this is the translated IP address of the device that was the network traffic's source.
Translated Port	sourceTranslatedPort	Integer	1	If network address translation is an issue, this is the translated source port associated with the attack.

Source Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Source Group Field Description
Translated Zone	sourceTranslatedZone	Zone	1	If network address translation is an issue, this is the network zone associated with the translated IP address of the device that was the network traffic's source.
Translated Zone External ID	sourceTranslatedZoneExternalID	String	1	Returns the external ID for this reference.
Translated Zone ID	sourceTranslatedZoneID	String	1	Returns the ID for the resource in this resource reference.
Translated Zone Name	sourceTranslatedZoneName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.
Translated Zone Reference ID	sourceTranslatedZoneReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Translated Zone Resource	sourceTranslatedZoneResource	Resource	1	Locates the resource described by this reference.
Translated Zone URI	sourceTranslatedZoneURI	String	1	Returns the URI for this reference.
User ID	sourceUserId	String	2	The OS- or application-based identifier associated with the user at the network traffic's source.
User Name	sourceUserName	String	2	The OS- or application-based name associated with the user at the network traffic's source.
User Privileges	sourceUserPrivileges	String	2	The privileges afforded the user at the network traffic's source.
Zone	sourceZone	Zone	1	The network zone where the source device resides.
Zone External ID	sourceZoneExternalID	String	1	Returns the external ID for this reference.
Zone ID	sourceZoneID	String	1	Returns the ID for the resource in this resource reference.

Source Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Source Group Field Description
Zone Name	sourceZoneName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.
Zone Reference ID	sourceZoneReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Zone Resource	sourceZoneResource	Resource	1	Locates the resource described by this reference.
Zone URI	sourceZoneURI	String	1	Returns the URI for this reference.

Target Group

Target Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Target Group Field Description
Address	targetAddress	IP address	1	The IP address of the device hosting the attacker. Note: For older versions of ESM, see "Device Custom Group" on page 929 for IPv6 address data type.
Asset ID	targetAssetId	Resource	2	The asset that represents the attacked device's host.
Asset Name	targetAssetName	String	2	The name of the device.
Asset Resource	targetAssetResource	Resource	2	See the common set of resource attributes.
DNS Domain	targetDnsDomain	String	2	The Domain Name Service domain name associated with the attacked device.
FQDN	targetFqdn	String	2	The fully qualified domain name associated with the attacked device.

Note about Geo fields:

Not all IPv6 addresses are mapped to Geo fields described below. In such cases, the corresponding Geo information for the IPv6 address will be blank on the Viewer panel.

Target Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Target Group Field Description
Geo	targetGeo	GeoDescriptor	1	The geographical information
Geo Country Code	targetGeoCountryCode	String	1	The identifier for the national-political state in which a device resides.
Geo Country Flag URL	targetGeoCountryFlagUrl	String	1	The URL of an image of the flag of the national-political state in which the device resides.
Geo Country Name	targetGeoCountryName	String	1	The name of the national-political state where a device resides.
Geo Descriptor ID	targetGeoDescriptorId	ID	1	The internal ID of the geographical reference.
Geo Latitude	targetGeoLatitude	Double	1	The latitude of a device.
Geo Location Info	targetGeoLocationInfo	String	1	Other, free-form text information about the device's location.
Geo Longitude	targetGeoLongitude	Double	1	The Longitude of a device.
Geo Postal Code	targetGeoPostalCode	String	1	The postal code of the device's location, as assigned by the national-political state where it resides.
Geo Region Code	targetGeoRegionCode	String	1	The identifier of the sub-region of the national-political state where a device resides. The style of the identifier varies with the host country.
Host Name	targetHostName	String	2	The name of the attacked device
MAC Address	targetMacAddress	MAC address	2	The MAC address associated with the target of the attack (which may or may not be the MAC address of the host device).
NT Domain	targetNtDomain	String	2	The Windows NT domain associated with the attacked device.
Port	targetPort	Integer	1	The network port associated with the target of the attack.

Target Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Target Group Field Description
Process ID	targetProcessId	Integer	2	The ID of the process associated with the attack's target.
Process Name	targetProcessName	String	2	The name of the process associated with the attack's target.
Service Name	targetServiceName	String	2	The name of service associated with the attack's target.
Translated Address	targetTranslatedAddress	IP address	1	If network address translation is an issue, this is the translated IP address of the attacked device.
Translated Port	targetTranslatedPort	Integer	1	If network address translation is an issue, this is the translated port associated with the attack.
Translated Zone	targetTranslatedZone	Zone	1	If network address translation is an issue, this is the network zone associated with the translated IP address of the targeted device.
Translated Zone External ID	targetTranslatedZone ExternalID	String	1	Returns the external ID for this reference.
Translated Zone ID	targetTranslatedZoneID	String	1	Returns the ID for the resource in this resource reference.
Translated Zone Name	targetTranslatedZoneName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.
Translated Zone Reference ID	targetTranslatedZone ReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Translated Zone Resource	targetTranslatedZone Resource	Resource	1	Locates the resource described by this reference.
Translated Zone URI	targetTranslatedZoneURI	String	1	Returns the URI for this reference.
User ID	targetUserId	String	2	The OS- or application-based identifier associated with the attacker, at the target of the attack.

Target Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Target Group Field Description
User Name	targetUserName	String	2	The OS- or application-based name associated with the attacker, at the target of the attack.
User Privileges	targetUserPrivileges	String	2	The privileges afforded the attacker, at the target of the attack.
Zone	targetZone	Zone	1	The network zone in which the attacked device resides.
Zone External ID	targetZoneExternalID	String	1	Returns the external ID for this reference.
Zone ID	targetZoneID	String	1	Returns the ID for the resource in this resource reference.
Zone Name	targetZoneName	String	1	Returns the name from the URI, which is always assumed to be the last field of the URI.
Zone Reference ID	targetZoneReferenceID	ID	1	Returns the unique descriptor ID for this reference. This is populated only if this reference has been stored and uniquely identified in the database.
Zone Resource	targetZoneResource	Resource	1	Locates the resource described by this reference.
Zone URI	targetZoneURI	String	1	Returns the URI for this reference.

Threat Group

Threat Group Data Fields

Label	Script Alias	Data Type	Default Turbo Level	Threat Group Field Description
Asset Criticality	assetCriticality	Integer	2	The relative measure of the importance of the targeted device, on a scale of 0 to 10.
Model Confidence	modelConfidence	Integer	2	The relative measure of ArcSight's confidence in its model of the attacked device, on a scale of 0 to 10.

Threat Group Data Fields, continued

Label	Script Alias	Data Type	Default Turbo Level	Threat Group Field Description
Priority	priority	Integer	1	The relative measure of importance of investigating this event on a scale of 0 to 10. This field incorporates Model Confidence.
Relevance	relevance	Integer	2	The relative measure of likelihood that this event succeeded, on a scale of 0 to 10.
Severity	severity	Integer	2	The relative measure of possible damage to network security represented by the event on a scale of 0 to 10. It may be noted that event severity is supplied by the device; connector severity is supplied by the SmartConnector; and attack severity is supplied by the threat evaluation process.

Resource Attributes

Resource Attributes Data Fields

Attribute Suffix	Description
External ID	The user-defined identifier associated with a configuration resource.
ID	The internal identifier associated with a resource (a UUID).
Reference ID	The internal identifier associated with the resource reference (an integer).
Type Name	The type of configuration resource.
URI	The URI associated with the resource (for example, /All Users/Administrators/Mlow).

Data Monitors

Data monitors are views within [Dashboards](#) that can be configured to report on events, [Filters](#), [Rules](#), and other areas that are of particular interest to you. Data monitors can be arranged on dashboards in numerous viewing layouts. Data monitors collect summary information on top events, most recent event activity, partial rule occurrences, hourly event counts, or event averages.

Data Monitors on Dashboards:

Once data monitors are created, they can be used to display information on dashboards. You can add one or more data monitors to the same dashboard to create a collection of different "instrument panel" monitors appearing in the Dashboard display in the Viewer panel. Both the

data monitors themselves and dashboards on which they are published can be shared among multiple Console users.

Permissions on Data Monitors:

Data monitors display only those events for which you have permission. In addition, if you do not have access to a data monitor, the data monitor does not function. Administrators can limit visibility of or control access to dashboards and data monitors by changing access control lists (ACLs) as needed. For more about this, see ["Managing Permissions" on page 107](#) and ["Controlling Who Has Permissions to Deploy Data Monitors" on page 116](#).

Data Monitor Types:

The ArcSight Console offers several predefined types to choose from when creating a new data monitor. The following topics describe the parameter entries and other options you can specify for each supported data monitor type.

You specify the Data Monitor type when you create a data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#). Also, the data monitors provided with ArcSight are examples of these various types of data monitors.

Asset Category Count Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

This data monitor enumerates the number of real-time hits (events) that occur per asset category, by priority, within a time interval.

Asset Category Count Data Monitor

Parameter	Description
Data Monitor Name	A unique name for the monitor.
Enable Data Monitor	Turn on the monitor and collect data from the Manager. If cleared, the monitor does not display data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Restrict by Filter	Choose a filter resource with which to restrict the events that can affect the asset categories.
Availability Interval	Set the number of seconds to use as the interval between monitor updates.

Asset Category Count Data Monitor, continued

Parameter	Description
Root Asset Category Group	Choose an asset-category resource group to monitor.
Levels	Set the number of resource hierarchy levels below the chosen Root Asset Group to monitor. A value of 1 monitors only the next level down. A value of -1 monitors all levels.
Aggregation	Turn on (True) or off (False) the ability to aggregate all hits to the asset group URI, including those above the leaf level, to reveal disparities or unanticipated counts that may merit drilling down.
Show Root URI	Choose whether to display (True) or not display (False) the complete URI for affected asset categories.
Show Root Series	Specify whether to include (True) or not include (False) the root series. This is used to select how many levels down in the hierarchy to include in the data monitor display. Using a combination of this, Show Root URI, Aggregation, and Levels, you can slice out single levels in the display.

Event Correlation Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

This data monitor provides flow-volume level correlation between two different event streams. The data monitor specifies two filters to identify two sub-streams of events within the overall stream of events coming into Manager. It then reports how closely the volume of events in the two streams correlate, that is, when the volume of events in Stream 1 decreases, does the volume in Stream 2 increase, decrease, or just change with no relation to the changes in Stream 1? For example, if a network intrusion detection system (NIDS) were deployed in front of several web servers in a cluster, one might expect that the flow of reported events from each NIDS would be roughly equivalent. If the event flow from one of the NIDS suddenly rose or fell out of sync with the other NIDS, then it might indicate a possible problem.

Event Correlation Data Monitor

Parameter	Description
Data Monitor Name	A unique name for the monitor.
Enable Data Monitor	Enable the data monitor and collect data from the Manager. If cleared, the associated viewer configuration does not display any data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278.
Restrict by Filter	Choose a filter resource with which to restrict the events that can affect the asset categories.
Availability Interval	Set the number of seconds to use as the interval between monitor updates.
Select Field Set	Specify a field set for use in data monitor drill-downs. When this data monitor is displayed, the user can double-click on a chart area or table row that represents an event to bring up a drill-down channel for that event. The field set specified here determines the columns (fields) shown in the drill-down channel. See also "Monitoring Dashboards" on page 252 for information on data monitor drill-downs.
Filter 1	Select a filter for the first event flow.
Filter 2	Select a filter for the second event flow.
Restrict by Filter	Choose to restrict the data monitor to a particular filter. When restricting by filter, you focus on a filter that is of particular interest to you and also reduce the number of events the data monitor retrieves.
Sampling Interval	Enter the interval (in seconds) for performing correlation calculations.
Number of Samples	Number of samples to keep in memory to perform calculations.
Availability Interval	Set the number of seconds to use as the interval between monitor updates.
Alarm Condition	Condition on which to fire an alarm, for example: <code>c > 90 && x > 0 && y > 0</code>. In this example, <code>c</code> represents the correlation count from -100 to + 100, <code>x</code> and <code>y</code> represent the actual count of events. See "Data Monitor Expressions" on page 993 for more information about the operators and functions supported in this and similar data monitor parameters that accept conditional expressions.
Maximum Alarm Frequency	Minimum time (in seconds) to wait before sending alarms for the same group.

How correlation is calculated

The event correlation data monitor applies covariance and correlation calculations to describe how two variables are related.

Covariance is calculated by the following formula:

$$COV(x,y) = \frac{\sum_{i=1}^n (x_i - \bar{x})(y_i - \bar{y})}{n - 1}$$

where:

x is the independent variable

y is the dependent variable

$\bar{x}$ is the mean of the independent variable x

$\bar{y}$ is the mean of the dependent variable y

Based on the covariance, correlation is then calculated by the following formula:

$$r_{(x,y)} = \frac{COV(x,y)}{s_x s_y}$$

where

$r_{(x,y)}$ is the correlation of variables x and y

$COV(x,y)$ is the covariance of variables x and y

s_x is the sample standard deviation of the random variable x

s_y is the sample standard deviation of the random variable y

Correlation standardizes the measure of interdependence between two variables and, consequently, tells you how closely the two variables move. The correlation measurement, called a correlation coefficient, will always take on a value between 1 and – 1:

- *If the correlation coefficient is 1*, the variables have a perfect positive correlation. This means that if one variable moves a given amount, the second moves proportionally in the same direction. A positive correlation coefficient less than one indicates a less than perfect positive correlation, with the strength of the correlation growing as the number approaches one.
- *If correlation coefficient is 0*, no relationship exists between the variables. If one variable moves, you can make no predictions about the movement of the other variable; they are uncorrelated.

- *If correlation coefficient is -1* , the variables are perfectly negatively correlated (or inversely correlated) and move in opposition to each other. If one variable increases, the other variable decreases proportionally. A negative correlation coefficient greater than -1 indicates a less than perfect negative correlation, with the strength of the correlation growing as the number approaches -1 .

The data monitor sampler takes all samples in memory and continually calculates correlation values using this formula. As an example, you could define an event correlation data monitor that displays a correlation between the number of times a network is being reconnoitered, and if that is related to the number of attacks that the network is receiving.

Event Graph Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

This data monitor draws real-time diagrams of selected event activity. In effect, it does automatically and in real-time what you can do manually, as described in ["Graphing Attacks" on page 287](#).



Note: The Zoom In, Zoom Out, and Fit Content options do not work with event graphs created after version 7.0. Scrolling works but can cause issues with large filter sets. To use these options with event graphs, you must enable the **Use classic charts** option in **Global Preferences**.

Event Graph Data Monitor

Parameter	Description
Data Monitor Name	A unique name for the monitor.
Enable Data Monitor	Select this check box to "switch on" the monitor and collect data from the Manager. If cleared, the monitor is "off" and displays no data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Restrict by Filter	Choose a filter resource with which to restrict the events that the graphic includes.
Availability Interval	Set the number of seconds to use as the interval between monitor updates.

Event Graph Data Monitor, continued

Parameter	Description
Select Field Set	Specify a field set for use in data monitor drill-downs. When this data monitor is displayed, the user can double-click on a chart area or table row that represents an event to bring up a drill-down channel for that event. The field set specified here determines the columns (fields) shown in the drill-down channel. See also "Monitoring Dashboards" on page 252 for information on data monitor drill-downs.
Show Event Nodes	Choose a basis for visually expanding or aggregating event nodes, relative to their source and target node instances. See "Configuring Event Graphs" on page 57 for the option details.
Max Event Count	Set the greatest number of most-recent events the graphic can show.
Show Source/Target Nodes as	When one source-event target chains to another, you can choose to graph a source/target IP address as a single (simple) node, or to graph both the source and target instances of such an IP address (distinct).
Source Node Identifier	Choose an event attribute to use as the identifier for source nodes. The default attribute is Source Address. Note that while all attributes are available, not all are appropriate choices for this purpose.
Event Node Identifier	The fields that are available to use to uniquely identify the event type in a transaction.
Target Node Identifier	Choose an event attribute to use as the identifier for target nodes. The default attribute is Target Address. Note that while all attributes are available, not all are appropriate choices for this purpose.

Geographic Event Graph Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

This data monitor draws a real-time geographic map of selected events. In effect, it does automatically and in real-time what you can do manually, as described in ["Graphing Attacks" on page 287](#).

Geographic Event Data Monitor

Parameter	Description
Data Monitor Name	A unique name for the monitor.
Enable Data Monitor	Select this check box to "switch on" the monitor and collect data from the Manager. If cleared, the monitor is "off" and displays no data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Restrict by Filter	Choose a filter resource with which to restrict the events that can affect the graphic. Filtering reduces the number of events the data monitor has to process. From the drop-down menu, double-click a filter or accept the default to receive all events.
Availability Interval	Sets the number of seconds to use as the interval between data monitor updates.
Select Field Set	Specify a field set for use in data monitor drill-downs. When this data monitor is displayed, the user can double-click on a chart area or table row that represents an event to bring up a drill-down channel for that event. The field set specified here determines the columns (fields) shown in the drill-down channel. (See "Monitoring Dashboards" on page 252 for information on data monitor drill-downs.)
Max Event Count	Set the greatest number of most-recent events the map can show.

Hierarchy Map Data Monitor

The data monitor type is chosen when you create a new data monitor. (For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).) This data monitor draws an image made up of proportionally sized panels where each panel represents a group of events selected by group fields selected in the source node identifier. A source-node criteria could be a combination of fields.

Related topics:

- ["Hierarchy Map Features" on the next page](#)
- ["Use Cases" on the next page](#)
- ["Defining a Hierarchy Map Data Monitor" on page 968](#)
- ["Adding Variables" on page 969](#)
- ["Specifying the Source Node Identifiers" on page 970](#)
- ["Specifying Group Attributes" on page 971](#)
- ["Hierarchy Map Display and Visualization Controls" on page 972](#)

Hierarchy Map Features

The Hierarchy Map data monitor includes the following features.

- The data monitor shows the complete hierarchy, with the hierarchy path built not just by using the delimiter within a field value but also across different field values. (Previous versions of the data monitor did not show the complete hierarchy.)
- *Group By* fields provide options to specify a list of delimiters for use by each selected Group by field. By default, no delimiters are used, if no delimiters are specified then the whole field is taken as a single level for hierarchy. (Previous versions built the hierarchy path within a field value based on only one type of separator, a forward slash, which did not support fields that use other separators like a backward slash, “\”, or a dot, “.”)
Group By fields also provide an option to set the maximum depth level of hierarchy within a field. The default depth level is equal to the number of delimiters in the field. Entering 0 for this option signifies no depth level for the selected field, effectively defining the field as a single-level hierarchy.
- A list of *Group Attributes* can be specified as a drill-down display to show when a user drills down into a group. For each attribute, the user can select a field and a function (max, min, count, average, count unique) on that field value.
- Enhanced visualization tools for *label*, *size by*, and *color by* provide fine-grained control of hierarchy map display with regard to Group By and Group Attributes fields and values.

Use Cases

Following is a list of example use cases for which the Hierarchy Map data monitor is a useful monitoring tool.

- Display the number of matches for all the rules within a given time frame, with the hierarchy groups based on the File path field of the rule audit events. The value is the count of the events for each group. The goal would be to show which rules fired the most in a given time frame.
- Show table space usage of correlation resources, particularly session lists and active lists.
- Show memory usage for correlation resources, particularly session lists and active lists.
- Show assets hierarchy by networks, zones and subnets. Within subnets, the assets can be sub-divided into asset ranges.
- Show assets hierarchy divided by the location of assets, where the value on the map is the count of the events targeting those assets.
- Show assets hierarchy divided by the location of assets, where the value on the map is the count of the assets within those locations.

- Monitor resource distribution; that is, how many rules, reports, data monitors and so on are being used in the system, where the count is system storage space.
- Display events by device to show how many events are generated from each device in a given time frame (for example, the past two days).
- Show assets by the number of attacks each receives, to determine which assets are the most vulnerable.

Defining a Hierarchy Map Data Monitor

1. Follow instructions in ["Creating a Data Monitor" on page 266](#) to add a data monitor.
2. In the Data Monitor editor, select **Hierarchy Map** as the Data Monitor Type .
3. Refer to the following table to set the attributes.

Hierarchy Map Data Monitor Attributes

Parameter	Description
Data Monitor Name	A unique name for the monitor.
Enable Data Monitor	Select this check box to "switch on" the monitor and collect data from the Manager. If cleared, the monitor is "off" and displays no data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Restrict by Filter	Choose a filter resource with which to restrict the events that can affect the graphic.
Availability Interval	Set the number of seconds to use as the interval between monitor updates.

Hierarchy Map Data Monitor Attributes, continued

Parameter	Description
Select Field Set	Specify a field set for use in data monitor drill-downs. When this data monitor is displayed, the user can double-click on a chart area or table row that represents an event to bring up a drill-down channel for that event. The field set specified here determines the columns (fields) shown in the drill-down channel. (See "Monitoring Dashboards" on page 252 for information on data monitor drill-downs.)
Source Node Identifier	This is a <i>group by</i> identifier. Blocks in the hierarchy map represents events or objects that have matching values for all fields chosen here. Also, identifiers specified here are available as Label By, Size By, and Color By choices on the displayed data monitor. Choose one or more event attributes by which to group events. The default attribute is Category Behavior, but you can include multiple attributes. For example, if you select only Category Behavior for this field, events are grouped by category behavior (for example, all events with a category behavior value of /Access are shown in one block, all events with a category behavior of /Authentication/Verify in another block, and so on). If you select more than one source node identifier, each block in the hierarchy map represents events or objects that have the same values for all identifiers. For example, if you select Category Behavior and Event Name as source node identifiers, then each block in the map represents events of the same behavior and event name. See "Specifying the Source Node Identifiers" on the next page for more information.
Group Attributes	You can specify one or more <i>group attributes</i> for fields with numerical values (for example, calculate the maximum priority of all events in a field group). The attributes you specify here are shown as drill-down tooltips when you mouse over a field on a hierarchy map display. You can add these attributes by specifying a label, a field, and a function to apply to the field. The functions can be applied on numeric fields only. See "Specifying Group Attributes" on page 971 for more details. Also, group attributes specified here are available as Label By, Size By, and Color By choices on the displayed data monitor.



Tip: If data monitor attributes are changed (edited) while a user is viewing the data monitor in a dashboard, the current data is flushed and the map defaults to red until new data arrives and the map display is redrawn.

Adding Variables

To add a variable to the Hierarchy Map data monitor:


1. On the data monitor editor, click the **Variables** tab.
2. For more on using variables in resources, see ["Variables" on page 1064](#).

Specifying the Source Node Identifiers

Source node identifiers are “group by” attributes. For example, if you select only Category Behavior for this field, events are grouped by category behavior. Each block in the hierarchy map represents a different type of category behavior (for example, /Authentication, /Authentication/Verify, /Execute Response/Informational, and so forth). If you select both Category Behavior and Target Address here, each block in the hierarchy map represents events with the same category behavior on the same target system (IP address or host name).

To specify one or more Source Node Identifier (Group By) fields:

1. Click in the **Source Node Identifier** field, then click the button  to open the Field Selector dialog.
2. Specify the fields by which you want to group events or objects by clicking Available Fields checkboxes, which adds them to “Fields to Show”. .
3. Click up/down arrows to re-order fields.
4. To remove a field, select it under Fields to Show and click the delete button .



For example, we can group by Category Behavior, Category Significance, and Target Address, which provide meaningful groups (events with the same category behavior, significance, and target address), and give us some interesting label, size, and color display options for mapping significant events and targeted systems on the data monitor.

Hierarchy Levels and Group Delimiters

You can specify how many levels of hierarchy you want to display for a field group by specifying one or more (a group of) delimiters and the maximum depth of hierarchy to display. For example, if you have a field value, `http://www.foo.com`, for which you have specified the depth level (Max Depth) as 2 with delimiters set to a group (consisting of ://.), you see:

- First level: http://
- Second level: http://www.foo.com

For the same example, if you set the Max Depth to 3, you get:

- First level: http://
- Second level: http://www
- Third level: http://www.foo.com

To select a field to display and set its hierarchy depth level:

1. Open the Hierarchy Map Field Selector dialog by clicking the browse button  that is displayed when you click in the Source Node Identifier field.
2. To add a field, check (click) the check box next to the field in the **Available Fields** scroll box. As you select a field, it is displayed in the Fields column in the “Fields to Show” table on the right side of the dialog.
3. Double-click the **Delimiter** column for the field you just selected and enter one or more delimiters based on which you want to show the hierarchy depth.
By default, a forward slash (/) is set as the delimiter. To set a single level of hierarchy, delete the “/” and do not specify any delimiters. Also, set the **Max Depth** (as explained in the next step) to zero for that field. If you set a comma (,) as a delimiter, the hierarchy in the panel displays a backslash (\).
4. To specify the depth of the field hierarchy within a field, double-click the **Max Depth** cell for the field.



Note: Negative integers are not allowed. If you enter a negative integer, it defaults to -1 which represents a depth level equal to the number of delimiters in the field.

If you leave this field blank, it defaults to a depth level equal to the number of delimiters in the field and -1 is displayed in the Max Depth column.

To display the whole field as a single level of hierarchy, set the **Max Depth** value to 0.

Specifying Group Attributes

Optionally, you can specify group attributes, which are functions on numerical fields. These attributes are shown as mouse-over tooltips on groups (blocks) on a displayed hierarchy map. They are also available as label, size, and color options.


For each attribute you want to add, provide a label, a function, and a field to which to apply the function. This can be done on numeric fields only. For example, if you add an Event Count label, select the Sum function, and apply this to the aggregatedEventCount field, the function finds the sum of event count value.

To add group attributes:

1. Click the **Group Attributes** cell. A browse button  is displayed.
2. Click the browse button. The Group Attributes dialog opens.
3. Click the **Label** column and enter a name for the attribute you want to create. You can add multiple labels.
4. Click the **Function** column for a label and select a function to be applied to a field that you select in the next step. You can set a function for a numeric field only.
5. Click the **Field** column against a label and select a field to which to apply the function.

For example, we'll create two labels, Event Count and Priority, and map them as follows.

Label	Function	Field
Event Count	Sum	assetCriticality
Priority	Max	priority



On the displayed map, the mouse-over tooltip on each block (group) shows both the event count and the highest priority events included in that block. Also, specified group attributes (Event Count and Priority, in this case) are available as Label By, Size By, and Color By options on the data monitor.


Hierarchy Map Display and Visualization Controls

After you create a Hierarchy Map Data Monitor, add it to a dashboard to display it so you can make further adjustments to the display.

Map Display and An Example

If no dashboards are displayed in the Viewer, simply right-click the Data Monitor you created, and select **Add to Dashboard As > Area Map**. This creates a new, untitled dashboard and add the data monitor to it. (You can also add it to an existing dashboard.)

The hierarchy map shown below is an example of the data monitor displayed on a dashboard.



You can choose **Hierarchy Map** as the Data Monitor type when you create a new Data Monitor. To display the data monitor, add it to a dashboard.



Tip: More reminders on working with the Hierarchy Map data monitor

- Before you can edit the visualization controls on the Hierarchy Map data monitor, add the data monitor to a dashboard and display the dashboard, as described in the beginning of this topic.
- If data monitor attributes are changed (edited) while another user is viewing that data monitor in a dashboard, the current data is flushed and the map defaults to red until new data arrives and the map display is redrawn.

The **example** data monitor above shows events grouped by category behavior, significance and target address. The labels show category behavior, and the blocks are sized by priority and colored by significance. Mouse-over tooltips show event count and priority for each group. Note that we can change the display on-the-fly by choosing a different label, size, and color options. For example, instead of coloring the blocks by Category Significance, we could color by Target Address. Or, instead of labeling by Category Behavior, we could label by Target Address. In this way, we can get quick, real-time, graphical overviews of network activity and adjust options to emphasize different details.

Labels, Size, and Color Controls

The visualization controls for Hierarchy Map Data Monitors are **Label By**, **Size By**, and **Color By** controls. (You might need to float the Viewer panel and expand the floating Viewer to see

these controls. See ["Changing the Console Display" on page 47.](#))

- **Label By** - Select a label. The value of the label you select is displayed on each block.
 - The **default** for Label By is all the fields specified for the source node identifier and the event count for that grouping. This shows as "Default" in the field. (The values available for use in the Label By field come from the attributes defined for Source Node Identifier and Group Attributes fields on the data monitor Editor. See ["Specifying the Source Node Identifiers" on page 970](#) and ["Specifying Group Attributes" on page 971](#) for more information.)
 - If Label by is set to something other than the default, the last (bottom-most) field value in the hierarchy does not show on the map because the custom Label by setting overwrites it. However, data for all fields, including the last field, is always taken into account on the map.

Use the default Label By option to show/visualize the complete hierarchy, including the last field value.

- **Size By** - Select an identifier or attribute by which you want to size the blocks. Once you select the Size By attribute, the blocks are resized proportionate to the value selected. Only attributes that have numeric values are available, because you cannot size a block based on a non-numeric value.


The values available for use in the Size By field come from the attributes defined for the Group Attributes field on the data monitor Editor. See ["Specifying Group Attributes" on page 971](#) for more information.

- **Color By**- Select identifier or attribute by which you want to color the blocks.

The values available for use in the Color By field come from the attributes defined for Source Node Identifier and Group Attributes fields on the data monitor Editor. See ["Specifying the Source Node Identifiers" on page 970](#) and ["Specifying Group Attributes" on page 971](#) for more information.

If you select a non-numeric field, you can change the color for any discrete value. If you select a numeric field, you get the option to select either a color for a discrete value or a color for a range of values. (For more on this option, see [Selecting Colors for the Blocks.](#))

After you select label, size, and color values, be sure to save the dashboard. The next time you open the dashboard, the attributes you saved are applied to the next set of data.



Format controls for the Hierarchy Map are available as drop-down menus on the map display in a data monitor.



Note: After an edit of tree map attributes, there might be a time lag before there is a visual indication of the updates. You can force a redraw of the tree map by dragging the slider to resize the panel that contains the map.

Selecting Colors for the Blocks

You can color the blocks by selecting any of the Source Node Identifiers or Group Attributes that are displayed in the **Color By** drop down menu. For example, if you select Priority in the **Color By** menu, then all blocks that have the same priority are displayed in the same color, such as all blocks with priority 1 may be displayed in red and all blocks with priority 2 may be displayed in blue, and so on.

If the Color By attribute you select is discrete but non-numeric, you can define the colors for each value of the attribute. For attributes that have numeric values, you can individually assign a color per attribute value or specify a range and assign a color for that range. However, if the Color By attribute is Priority, you cannot specify a range. This is because there are already predefined colors for each level of priority. You can change a predefined color to a color of your choice for each priority level.

Below the Label By, Size By and Color By fields, is the Color Chooser box. This box displays all the values for the **Color By** group/field that you select. To individually assign a color for an attribute:

1. Click the **Discrete** radio button (This button is visible only if you selected a numeric Color By attribute).

2. Double-click a value button to open the **Color Chooser** dialog.
3. Select a color that you want to display for all the boxes for which that value is applicable.
4. Click **OK**.

All the boxes that have that value is displayed in the new color.

You can set a threshold for the maximum number of discrete values for which you can set a color. Set the `console.ui.hmDataMonitor.discrete.threshold` property in the `console.defaults.properties` file. If the number of discrete values exceeds this threshold, for all values that cross the threshold, the color is set to white.

To assign a color for a range of values (for numeric fields only):

1. Click the **Range** radio button.
2. Click **Add** button to set a range and a color for that range. The Add a color mapping dialog opens.
3. Select a value from the **Min Attribute Value** and **Max Attribute Value** menus to set the range.

For example, if you want to set a range for Priority that falls in 3-to-6 range, select 3 from the Min Attribute Value menu and 6 from the Max Attribute Value menu.

4. Click the Color Chooser button to open the color chooser.
5. Select a color by clicking it and click **OK**. The color you choose is used to display all values falling in that range. In our range example in step 3, all blocks that display priority of 3, 4, and 5 have the color you just chose for the 3-6 range.



Tip: If new data comes in after you change the color mapping but before you save the new mapping, you get a dialog asking you whether you want to save the changed mapping. If you select **Yes**, the Data Monitor is not refreshed with new data until you save the new mapping. When you save it, the new mapping is applied to the existing blocks and all future data displayed on the dashboard.

If you select **No**, the new color mapping is applied to the existing data on the dashboard, but is not saved in the database. So, as soon as new data arrives, the new color mapping is overwritten by the original color mapping that exists in the database.

Hourly Counts Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

The Hourly Counts Data Monitor displays the total count of events on an hourly basis along with their Priority. The hourly count for the first hour segment starts when you open the

dashboard. For example, if you open the dashboard at 2:25 PM, though the first time segment displays **14:00 - 15:00**, the count begins at **2:25 PM**.

Hourly Counts Data Monitor

Parameter	Description
Data Monitor Name	Enter a data monitor name. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Enable Data Monitor	Select the check box to enable the data monitor and collect data from the Manager. If not selected, the associated viewer configuration will not display any data.
Restrict by Filter	Choose a filter resource to restrict the data monitor's contents.
Availability Interval	Set the number of seconds to use as the interval between monitor updates.
Select Field Set	Specify a field set for use in data monitor drill-downs. When this data monitor is displayed, the user can double-click on a chart area or table row that represents an event to bring up a drill-down channel for that event. The field set specified here will determine the columns (fields) shown in the drill-down channel. (See "Monitoring Dashboards" on page 252 for information on data monitor drill-downs.)

As an example, you could design an Hourly Counts data monitor that displays hourly counts of data being collected, for example, the number of events that Manager receives.

Last N Events Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

The Last *N* Events data monitor orders events based on its configuration. In the Table Viewer, the monitor displays the most recent events by Priority, Event Name, Protocol, and Category. With the BarChartTable configuration, the order is by Priority and Event Name. The PieChart configuration is ordered by Priority.



Note: If your Last *N* Events data monitor includes a column that displays the annotation stage, that column is not updated when an analyst later changes an event's annotation stage. Data monitors are designed to display events as they flow in for the first time. For annotation updates, OpenText recommends using query viewers, which are configurable to re-query the database, then add the query viewer to the dashboard.

Last N Events Data Monitor

Parameter	Description
Data Monitor Name	Type a data monitor name.
Enable Data Monitor	Select the check box to enable the data monitor and collect data from the Manager. If not selected, the associated viewer configuration will not display any data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Availability Interval	Set the number of seconds to use as the interval between monitor updates.
Restrict by Filter	Choose a filter resource to use as an additional restriction on the events displayed.
Select Field Set	Specify a field set for use in data monitor drill-downs. When this data monitor is displayed, the user can double-click on a chart area or table row that represents an event to bring up a drill-down channel for that event. The field set specified here will determine the columns (fields) shown in the drill-down channel. (See "Monitoring Dashboards" on page 252 for information on data monitor drill-downs.)
# of Events	Specify how many events the data monitor displays.
Field Names	Choose field names to include in the data monitor display. By default, the data monitor includes EventName, EventCategory, ArcSight Severity, and Protocol fields. You can select additional fields or remove currently selected fields by Shift or Ctrl-clicking field names in the drop-down list.

As an example, you could design a Last N Events data monitor that displays the latest N events that meet the condition specified in the dashboard definition.

Last State Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

This monitor is somewhat different than others in that it provides an extra level of abstraction that you can use to simplify the information presented to operators. Sometimes called "indicator lights" or "heads-up displays," these monitors show graphics that translate more complex values into simple, rapidly observable results such as green/amber/red "signal lights" or checkmark/asterisk/exclamation point symbols. "Last State" data monitors could also be called "most recently known state" monitors.

Last State data monitors are built on the information collected by [Active Lists](#). The qualifying events in active lists are identified on the basis of selected key fields such as Source Zone and Source Address (see [Data Fields](#)).

Having focused on the events that apply, you then select a field to use as the basis of the values the indicators will simplify. For example, the Priority field has a range of values you could divide into sub-ranges that you choose to translate into good/okay/bad groups.

With a value-range and status-scheme decided, you can map the field values to the status names, and the status names to the visual indicators operators will see.

Last State Data Monitor Parameters

Last State Data Monitor

Parameter	Description
Data Monitor Name	A unique name for the monitor.
Enable Data Monitor	Select this check box to "switch on" the monitor and collect data from the Manager. If cleared, the monitor is "off" and displays no data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Restrict by Filter	Choose a filter resource to use as an additional restriction on the events summarized through the indicators, if necessary.
Availability Interval	Set the number of seconds to use as the interval between monitor updates.
Select Field Set	Specify a field set for use in data monitor drill-downs. When this data monitor is displayed, the user can double-click on a chart area or table row that represents an event to bring up a drill-down channel for that event. The field set specified here will determine the columns (fields) shown in the drill-down channel. (See "Monitoring Dashboards" on page 252 for information on data monitor drill-downs.)
Restrict by Active List	Choose an active list from the resource tree to use as the primary guide for event selection. The choices are limited to event-based active lists.
Key Fields	Choose the fields to use as identifiers for the indicators, and the order in which to display them.
Value Fields	Select the fields that provide the range of values to be mapped into indicators, and the order in which they are to be evaluated.
Max Number of Indicators	Set the greatest number of qualifying indicators the data monitor will show. If more indicators are generated, the displayed set will be the those with the most recent event traffic.

Last State Data Monitor, continued

Parameter	Description
Mapping	Use the Define Status Map dialog box in two steps: first, on the Statuses tab, to associate status Titles with Image graphics, then, on the Mapping tab, to associate Value items contained by the Value Fields with the Statuses titles just defined. Be sure to define and select one "catch all" status to react to values that may fall outside the range you set. In the Value field, associate only one value at a time with the Status values you've defined. For example, if the values 0, 1, and 2 should all be associated with a Status of "Okay," enter each digit separately and click Add.
Use as Timestamp	Choose whether to use the device's reported end-time or the Manager's receipt time as the definitive timestamp.
History Function	Use this option to add a Min or Max column to grid views that shows the minimum or maximum value for the indicator, over the most-recent time period specified by History Time Range.
History Time Range	Used with the History Function. The (most recent) period of time, in minutes, for which to retain minimum or maximum value information for an indicator. For example, a value of 60 could cause an indicator's Max column in a table to show its highest registered value over the previous hour.
Timeout	Used with the History Function. Sets the time limit, in seconds, after which the Min and Max column values are purged if not already updated.

Options for Table and Tile Views

In dashboards, you can see Last State data monitors as **Table** or **Tile** views. Click the **View as icon** () button at the lower-right corner to choose Table or Tile view.

The Table view will show more items than a *modified* Tile view. If the Tile view is customized to show results with particular values for key fields, it will show only a subset of the data monitor results.

A color chooser is available to apply to the **Table** view.


Table View (Color Chooser and Remove Entry)

Also in Table view, you can right-click an entry in a Last State data monitor and choose **Remove Entry**. However, keep in mind the data monitor's [Availability Interval](#) setting. Removal does not visibly take place until the next refresh, during which time a new instance of the entry could occur. Depending on the entry and the interval, a removed entry may appear to have remained.


Tile View (Customize View)

When in **Tile** view, you can use the **Customize** button () to change the way data is ordered in the tabular (tiled) presentation, and limit the view to a subset of data monitor results. The

customization choices are **by row-and-column** and **by cell**. Row-and-column is quicker to set up than cell because there are fewer adjustments, but cell does give you the option to set the contents of each tile in the data monitor. The custom settings here, in effect, filter your view based on values you specify for key fields. For example, if you are interested in monitoring state changes on only four systems identified by target host name and target address; you can customize the view on a dashboard to show only those four systems.



These tiled views are "fixed" meaning that the tiles in the array will hold their positions, relative to each other and to the dashboard.



The customize view option on the Last State Data Monitor tile view gives Console users a way to design their own views and then get quick visuals on a few key assets or attackers:

- Set up a custom, focused view of a few items that you are interested in (assets like servers that might be targets of attacks, suspicious nodes that might be attackers, and so forth). You do this by submitting values for the items you want to monitor (for example, host names and addresses, if those are your key fields).
- In a single glance, get information on state changes in different priority events (low, medium, high) in these few items. Since the positions of the items in your custom tile view does not change, and because you have limited the view to a few key items, you can get last state status with a quick glance.
- ArcSight Console users can customize dashboard views that use the same underlying Last State Data Monitor, but focus in on different items or assets (depending on how the data monitor custom view is set up on each dashboard). Fred could have a dashboard set up to monitor a few key servers based on host name and address, while Ethel could have a dashboard set up to monitor some firewalls, and both could be using the same underlying data monitor.



Tip: Notes on the Customize view option for Last State Data Monitor

- This is a dashboard-level customization that essentially filters the view based on values you provide for key fields. The key fields that show up in the Viewer Settings come from the key fields set in the data monitor.
- A priority mapping needs to be configured in the data monitor in order for the quick-glance tile view to provide useful last state information
- Customizations made to the Last State Data Monitor are saved to the dashboard that holds the data monitor, not to the data monitor itself. This allows the same data monitor to be used as a basis for different, customized views (dashboards) for quick-glance, priority state changes related to incoming events.
- The Customize view option applies only to the **Tile** view (not Table view). When you switch to the **Table** (), you can see results for other items in addition to the ones you “filtered” for on the customized Tile view.

Moving Average Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

The Moving Average data monitor displays the moving average of events by a selected data field. The display provides a running count of events within a specified time frame and generates an event when the moving average changes significantly.

If a Moving Average data monitor is configured to display multiple graphs simultaneously, you can open it using the Statistics Chart or Tile format options described in ["Creating or Editing a Dashboard" on page 255](#).

This data monitor calculates its statistics based on the number of requested samples. Until a full set of samples accumulate, the statistics approach their nominal value. This is indicated by appending /Partial to the event category if the values represent an incomplete sample. The purpose is to prevent false positives. This is most applicable to /DataMonitor/MovingAverage/Threshold/ events.

When either the Moving Average or Statistics data monitors gain or lose a value grouping during processing (for example, Priority), they issue an internal event. The data monitor's event categorization shows a Value/Add or Value/Remove suffix. This makes it possible to detect anomalous drops to zero, which can otherwise be missed if the monitor is removed because the discard threshold and a Threshold/Falling event could not be sent (due to exceeding the Maximum Alarm Frequency setting).

Both the Moving Average and Statistics data monitors have a **Stats Value Field**. When used, this attribute focuses the monitor's statistical analysis on the numeric value of a specified field rather than on the quantitative flow of events. Analyzing numeric fields within events enables a broad number of possibilities for status monitoring, especially with custom strings and ArcSight [Audit Events](#).

The **Value Calculation** field offers additional time-sensitive options for monitoring in second or minute increments. Monitoring per-second can catch abrupt spikes or drops; monitoring per-minute allows the same capability but may be more appropriate for larger integer values.

Moving Average Data Monitor

Parameter	Description
Data Monitor Name	Type a data monitor name.
Enable Data Monitor	Select the check box to enable the data monitor and collect data from the Manager. If not selected, the associated viewer configuration will not display any data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Restrict by Filter	Specifies whether to restrict the data monitor to a particular filter. When restricting by filter, you focus on a filter that is of particular interest to you and also reduce the number of events the data monitor retrieves. From the drop-down menu, double-click a filter or accept the default to receive all events.
Availability Interval	Set the number of seconds to use as the interval between monitor updates.

Moving Average Data Monitor, continued

Parameter	Description
Select Field Set	Specify a field set for use in data monitor drill-downs. When this data monitor is displayed, the user can double-click on a chart area or table row that represents an event to bring up a drill-down channel for that event. The field set specified here will determine the columns (fields) shown in the drill-down channel. (See "Monitoring Dashboards" on page 252 for information on data monitor drill-downs.)
Stats Value Field	Specify a particular numeric field within events to use for statistical evaluation, rather than the overall flow of events. For example, specifying the Priority field would focus the data monitor on changes to the value of the Priority field in events, instead of on changes to the number of events encountered. The default is Aggregated Event Count, which is the sum of all aggregated events. Tip: Events can be <i>aggregated</i> at the Connector on specified fields. This pares down the number of events of the same type that the Manager must process.
Value Calculation	Controls the way the time-based accumulation of values is evaluated against the number of events involved. The default is Sum of values, which is the sum of all Stats Value Field event values. Average value per event divides the value by the number of events in the unit. Average value per second divides the value by the number of seconds in the unit. Average value per minute divides the value by the number of minutes in the unit. For finer time-sensitive value calculations, also consider using the Number of Samples and Sampling Interval so results are neither too shallow or too acute to be meaningful.
Group By	Group by the specified field (for example, Priority)
Sorted By	Sort by the values found in fields or by the percentage of change in those values.
Alarm Change Threshold (%)	Specifies the moving average threshold, the percent change from the moving average, that will send a threshold exceeded event to the ArcSight Console. The threshold exceeded event is sent to the Console and can be used to create a rule. For more information on rules, see "Managing Rule Actions" on page 532 . Type in a percentage. The default is 50.
Number of Samples	Type the number of Sampling Intervals to use to calculate the moving average, in seconds. The most recently stored Sampling Intervals are used to calculate the moving average. For example, if five Number of Samples are used, the last five Sampling Intervals are used to calculate the moving average.
Number of Visible Groups	Set the number of rows of results to display in the data monitor for each combination of ordering fields specified in the Group By parameter.

Moving Average Data Monitor, continued

Parameter	Description
Sampling Interval	Type the time interval used to calculate the moving average, in seconds. For example, if the Sampling Interval is 5 minutes, the moving average is calculated every 5 minutes. The default is 300.
Group Discard Threshold	Specifies the minimum event counts needed to generate a threshold exceeded event. For example, event count could change from 1 to 2, a 100% change that results in a threshold exceeded event. To prevent these types of changes from generating a threshold exceeded event, specify the minimum event counts needed. If you want all events generated regardless of the event count, type 0.
Maximum Alarm Frequency	Minimum time (in seconds) to wait before sending alarms for the same group.

For example, you could design a Moving Average data monitor that displays the moving average of events on a per-source-address basis.

In ArcSight Reports/Custom Reports/Moving Average Report, you can specify the name of the dashboard as a parameter (same as the moving average event name), and specify the detect time range to report on.



Note: You can also have a rule trigger based on the moving average of events coming in, independent of defining reports based on moving average events.

Rules Partial Match Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

Displays rules that have partial matches and the total number of partial match events within a specified time frame. See also ["Automatically Disabled Rules " on page 1035](#).

Rules Partial Match Data Monitor

Parameter	Description
Data Monitor Name	Type a data monitor name.
Enable Data Monitor	Select the check box to enable the data monitor and collect data from the Manager. If not selected, the associated viewer configuration will not display any data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .

Rules Partial Match Data Monitor, continued

Parameter	Description
Window Size	Specifies the time interval used to report partial match counts, in seconds. For example, if using 1 hour as the Window Size, each window displays partial match counts in hour intervals. The default is 3600.
Number of Windows To Display	Type the number of Window Sizes to display. The default is 5.
Fixed or Sliding	Specifies when to begin the Window Size time interval. Choose Fixed to begin at time units, such as every hour, 1:00, 2:00, and so forth, or Sliding to begin at the current time and move backwards in Window Size time intervals. For example, if the window size is 10 minutes, and the current time is 1:15 PM and Fixed was selected, the window time frames would be 1:00 to 1:09 and 1:10 to 1:15. If Sliding was selected, window time frames would be 1:00 to 1:04 and 1:05 to 1:15.

For example, you could design a Rules Partial Match data monitor that displays all events that have partially matched and enabled real-time rule conditions, and are currently stored in memory.

Statistics Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

The Statistics Data Monitor provides a broader generalization of Moving Average data monitor functionality, except that it allows selection of other statistical methods in addition to Moving Average. Statistical methods include Average, Moving Average, Standard Deviation, Skew, and Kurtosis. These added capabilities could be used to detect anomalous behavior that could not be detected using moving average alone.

For example, monitoring the standard deviation of event data allows alarms to be triggered when there are sudden shifts in the rate of change of an event flow. This would allow alarms to be triggered when the protected network has been infected with a worm, but not when the network traffic rises due to normal use.

Both the Statistics and Moving Average data monitors have a **Stats Value Field**. When used, this attribute focuses the monitor's statistical analysis on the numeric value of a specified field rather than on the quantitative flow of events. Analyzing numeric fields within events enables a broad number of possibilities for status monitoring, especially with custom strings and [ArcSight Audit Events](#).

In dashboards, you can see Statistics data monitors as **Statistics Chart** or **Tile** views. Click the **View as icon** button () at the lower-right corner to choose. When in Tile view, you can use the **Customize** button () to change the way data is ordered in the tabular (tiled) presentation. The customization choices are **by row-and-column** and **by cell**. Row-and-column is quicker to

set up than cell because there are fewer adjustments, but cell does give you the option to set the contents of each tile in the data monitor.

When either the Moving Average or Statistics data monitors gain or lose a value grouping during processing (for example, Priority), they issue an internal event. The data monitor's event categorization shows a Value/Add or Value/Remove suffix. This makes it possible to detect anomalous drops to zero, which can otherwise be missed if the monitor is removed because the discard threshold and a Threshold/Falling event could not be sent (due to exceeding the Maximum Alarm Frequency setting).

These tiled views are "fixed," meaning that the tiles in the array will keep their positions, relative to each other and to the dashboard.

Statistics Data Monitor

Parameter	Description
Data Monitor Name	Enter a data monitor name.
Enable Data Monitor	Select the check box to enable the data monitor and collect data from the Manager. If not selected, the associated viewer configuration will not display any data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Restrict by Filter	Choose to restrict the data monitor to a particular filter. When restricting by filter, you focus on a filter that is of particular interest to you and also reduce the number of events the data monitor retrieves.
Availability Interval	Set the number of seconds to use as the interval between monitor updates.
Select Field Set	Specify a field set for use in data monitor drill-downs. When this data monitor is displayed, the user can double-click on a chart area or table row that represents an event to bring up a drill-down channel for that event. The field set specified here will determine the columns (fields) shown in the drill-down channel. (See "Monitoring Dashboards" on page 252 for information on data monitor drill-downs.)
Statistics Type	Choose the type of statistical calculation the data monitor will perform. The available types are Average, Identity, Kurtosis, Skew, Standard Deviation, and Variance.
Stats Value Field	Specify a particular numeric field within events to use for statistical evaluation, rather than the overall flow of events. For example, specifying the Priority field would focus the data monitor on changes to the value of the Priority field in events, instead of on changes to the number of events encountered.
Group By	Group by the specified field (for example, Name)
Sorted By	Choose to sort results by value, sample count, statistics, or triggering criteria.

Statistics Data Monitor, continued

Parameter	Description
Alarm Trigger Condition	Enter a conditional expression on which to trigger alarms. You can use any mathematical expression that employs these three variables, using <i>n</i> as the Number of Samples: c = The new sample ps = Statistics from previous <i>n</i> samples <i>excluding</i> c s = Statistics from the last <i>n</i> samples <i>including</i> c For example, the following expression would trigger when the current sample goes beyond 500: <pre>c >= 500</pre> An expression that triggers when the statistics reach 500 would be: <pre>s >= 500</pre> As a matter of interest, the Moving Average data monitor is in effect a special case of the Statistics data monitor, based on this expression: $s \neq 0 \ \&\& \ (abs((c - s)/s) * 100) > 50$ where 50 is the percent of change you specify in the Moving Average data monitor. See "Data Monitor Expressions" on page 993 for more information about the operators and functions supported in this and similar data monitor parameters that accept conditional expressions.
Number of Samples	Specify the number of most-recent Sampling Intervals to retain in memory and use to calculate event statistics. For example, if you set it to retain 5 sampling intervals, the last five periods (as specified in the Sampling Intervals attribute) are used to calculate the moving average.
# of Groups to Display	Set the number of rows of results to display in the data monitor for each combination of ordering fields specified in the Group By parameter.
Sampling Interval	Enter the time interval for recalculating event statistics, in seconds. For example, if the Sampling Interval is 5 minutes, the moving average is calculated every 5 minutes.
Group Discard Condition	Enter a condition (a filtering expression) by which to remove certain result rows from consideration in statistical calculations, based on the result ordering set in the Group By attribute. See "Data Monitor Expressions" on page 993 for more information about the operators and functions supported in this and similar data monitor parameters that accept conditional expressions.
Maximum Alarm Frequency	Minimum time (in seconds) to wait before sending alarms for the same group.

System Monitor Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

The System Data Monitor provides measurements based on Manager internal monitoring system Java classes and attributes. A number of system monitors that may be particularly useful to administrators are provided as predefined System Data Monitors that you can include in your dashboard displays to monitor system performance.

System Monitor Data Monitor

Parameter	Description
Data Monitor Name	Type a data monitor name
Enable Data Monitor	Select the check box to enable the data monitor and collect data from the Manager. If not selected, the associated viewer configuration will not display any data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Monitor Types	From the drop-down menu, select the name of ArcSight Java class for which you want to display attribute measurements, for example, Throughput meter or Status

System Monitor Attribute Data Monitor

The data monitor type is chosen when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

The System Monitor Attributes Data Monitor is similar to System Monitor, except that, rather than providing measurements for all attributes of a specified Java class, focuses on a single specific attribute of a given ArcSight Java class. (Used primarily for measurements on attributes that provide complex data structures.) A number of predefined system monitors are provided that you may want to include in your dashboard displays to monitor system performance.

System Monitor Attribute Data Monitor

Parameter	Description
Data Monitor Name	Type a data monitor name.
Enable Data Monitor	Select the check box to enable the data monitor and collect data from the Manager. If not selected, the associated viewer configuration will not display any data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .

System Monitor Attribute Data Monitor, continued

Parameter	Description
Monitor Types	From the drop-down menu, select the name of ArcSight Java class for which you want to display attribute measurements, for example, Throughput meter or Status.
Attribute Name	Specify the individual attribute of the specified ArcSight Java class for which you want to display information. You can obtain the names of specific attributes in a class by viewing the results of a System Monitor defined for that class.

Top Value Counts Data Monitor

This data monitor type is a selection when you create a new data monitor. For information on how to create a data monitor, see ["Creating a Data Monitor" on page 266](#).

The Top Value Counts data monitor displays top events by selected data fields, the total number of events, and the event Severity within the total number of events as defined by the filter (Restrict by Filter parameter). Data is displayed in Table and BarChartTable viewer configurations.

Top Value Counts uses an aggregation mechanism that precisely and predictably controls the time dimension of the data being evaluated. "Bucketized" means that the monitor evaluates a specific number of time-based event data units of a certain size (buckets). As time increments forward, the evaluation refreshes, using the most recent set of qualifying buckets.

The data monitor's latest bucket process live data. You should expect some delay ranging from milliseconds to seconds between the Manager's receipt of the event and when the event is processed by the data monitor. The latest bucket may therefore not have counted all the events up to the current millisecond. Eventually the count discrepancy is resolved and the bucket counts will be correct.

Top Value Counts Data Monitor

Parameter	Description
Data Monitor Name	Enter a data monitor name.
Enable Data Monitor	Select the check box to enable the data monitor and collect data from the Manager. If not selected, the associated viewer configuration will not display any data. Depending on the permissions associated with the user group to which you belong, you may or may not have an option to Enable (<i>deploy</i>) or disable (<i>un-deploy</i>) the data monitor. For more information, see "Enabling or Disabling a Data Monitor" on page 278 .
Restrict by Filter	Specify a filter to focus on events that are of particular interest and to reduce the number of events the data monitor processes. Use a filter when the number of possible Aggregate Field values can exceed the maximum for # of Distinct Events .

Top Value Counts Data Monitor, continued

Parameter	Description
Availability Interval	Sets the number of seconds to use as the interval between monitor updates.
Select Field Set	Specify a field set for use in data monitor drill-downs. When this data monitor is displayed, the user can double-click on a chart area or table row that represents an event to bring up a drill-down channel for that event. The field set specified here will determine the columns (fields) shown in the drill-down channel. (See "Monitoring Dashboards" on page 252 for information on data monitor drill-downs.)
Bucket Size in Seconds	The time dimension for individual event data units. A number of these units make up the value used in Number of Buckets . For example, you might use a value of 300 to create five-minute buckets. Bucket size and frequency (increasing freshness and resolution) does have a performance cost so it is wise to set buckets to run only as small and fast as actually necessary.
Number of Buckets	The overall time dimension to evaluate, expressed as the appropriate number of Bucket Size units. For example, to evaluate the most recent hour using five-minute buckets, you would enter 12 . Bucket size and frequency (increasing freshness and resolution) does have a performance cost so it is wise to set buckets to run only as small and fast as actually necessary.
Time Field	Choose the specific event timestamp to use to apply events to time buckets.
# Top Entries	The number of entries to show as "top" values.
# of Distinct Events	This value must equal or exceed the maximum number of values that the Aggregate Field can possibly have. The default is 1,000. The maximum is 10,000. This value controls the upper limit on the number of aggregate field values. If it is smaller than necessary, then when it encounters one more Aggregate Field value than allowed, the Data Monitor resets all the counters, clears the data, and starts over at zero. If you specify more than one Aggregate Field, the maximum number of possibilities is the product of the possible values of all fields. For example, if you are aggregating by users and zones in an environment with 200 users and 15 zones, the number of possibilities is $200 \times 15 = 3,000$. If the number of possibilities is larger than the maximum of 10,000, use a filter to reduce them.

Top Value Counts Data Monitor, continued

Parameter	Description
Aggregate Field	Specify one or more data fields to monitor. For more information, see "Data Fields" on page 910 . To monitor the top 10 source IP addresses, for example, select the Source Address data field from the drop-down menu. If you specify more than one field, the total number of possible combinations is the product of the number of possible values for each field you specify. Make sure that the # of Distinct Events field is large enough to accommodate this number.
Value Field	Specify what the data monitor will use when determining the top value counts: the number of matching events, or the sum of a particular data field value in all matching events. • To count events, leave this field empty. (This is equivalent to selecting the Aggregated Event Count field. When the Value Field is not specified, the data monitor uses the data field specified in the Aggregate Field to count events.) • To sum the values from a particular data field, use the data field selector for the "Value Field" attribute to select the desired field. In either case, counts from aggregated events will be properly adjusted.
Send Audit Events	Specify generation of audit events for this data monitor. By default, audit events are not generated. Refer to "Audit Events" on page 833 and look for the audit events under "Top Value Counts Data Monitor."

Troubleshooting

You might see warnings about the Top Value Counts data monitor type in the server logs, stating that internal data structures are being discarded to prevent overflow. "Data structures" in this warning refer to the counts being tracked. The events are not actually lost. This warning indicates the data monitor is using system resources but not providing any useful statistical data because the data monitor's conditions are poorly selected. When this warning appears, the problem continues until you fix your data monitor configuration.

Try these:

- Choose your event filter and aggregation settings carefully. For example, if you are counting top *N* source addresses from within your organization, you should not see these warnings. But if your top *N* source addresses are from a broader source like the internet, the top 1000 will be easily exceeded and the data will be flushed. This means counting will start all over repetitively.
- The data monitor will monitor up to 1000 distinct matches (can be increased to up to 10000) based on the event filter and aggregation settings and from that, determine the top *N* (10 by default). Increasing the # of Distinct Events default from 1000 to a maximum 10000 might help, but will cause the data monitor to consume more resources.

Data Monitor Expressions

Certain data monitor parameters can specify their own conditional expressions with which to flexibly define triggers or results. For example, you use these expressions in the Statistics data monitor's Alarm Trigger Condition and Group Discard Condition parameters to evaluate when to send an alarm or to remove result rows from statistical calculations.

The type of expression supported is a conventional infix mathematical expression with each basic expression separated by parentheses.

All common arithmetic operators are supported. Boolean operators are also fully supported and Boolean expressions evaluate as either 1 or 0 (true or false).

Supported Data Monitor Expression Operators

All common arithmetic operators are supported. Boolean operators are also fully supported and Boolean expressions evaluate as either **1** or **0** (true or false).

Data Monitor Expression Operators

Operator	Symbol	Operator	Symbol
Power	^	Less Than or Equal	<=
Boolean Not	!	More Than or Equal	>=
Unary Plus	+x	Less Than	<
Unary Minus	-x	Greater Than	>
Modulus	%	Not Equal	!=
Division	/	Equal	==
Multiplication	*	Boolean And	&&
Addition	+	Boolean Or	
Subtraction	-		

Supported Data Monitor Expression Functions

Data Monitor Expression Functions

Name	Function	Name	Function
Sine	sin()	Inverse Hyperbolic Cosine	acosh()
Cosine	cos()	Inverse Hyperbolic Tangent	atanh()
Tangent	tan()	Natural Logarithm	ln()

Data Monitor Expression Functions, continued

Name	Function	Name	Function
Arc Sine	asin()	Logarithm Base 10	log()
Arc Cosine	acos()	Angle	angle()
Arc Tangent	atan()	Absolute Value / Magnitude	abs()
Hyperbolic Sine	sinh()	Random Number (between 0 and 1)	rand()
Hyperbolic Cosine	cosh()	Modulus	mod()
Hyperbolic Tangent	tanh()	Square Root	sqrt()
Inverse Hyperbolic Sine	asinh()	Sum	sum()

Device

- See ["Assets" on page 830](#) for a discussion of network devices.
- See ["Device Group" on page 924](#) and ["Device Custom Group" on page 929](#) for information on device-related data fields.

Event Inspector

The Event Inspector is a tool for examining event details (see ["Events" on the next page](#) and ["Event Categorization" on page 997](#) for information about events). The Event Inspector is located in the ArcSight Console's Inspect/Edit panel. To open the Event Inspector, double-click an event line in a grid view. See ["Views" on page 1096](#).

There are two panels in the Event Inspector. The top panel displays selected events with associated rules. The events listed here have a set of right-click menu commands similar to those described in ["Using Active Channel Menu Commands" on page 241](#). The bottom panel displays event details for one or more events that have been selected from the top panel. If you select more than one event from the top panel, only their common values are displayed in the bottom panel.

The Event Inspector can display the chain of events that trigger a rule (see ["Rules" on page 1034](#)) and generate a correlation event. From the Event Inspector you can view each event and rule in the chain for details.

Depending on the information available for an event, you may also be able to review its business significance in the Impact Analysis tab or its actual content in the Payload tab.



Tip: Viewing global variables in the Event Inspector

When you view events in an active channel and open an event that contains a global variable field in the Event Inspector, you may need to refresh the Event Inspector view to see the global variable fields, because the Manager processes global variable data differently from regular event data.

- If your Hide Empty Rows icon  is toggled on (so that empty rows are not displayed), you may not see global variable field in the event inspector.
- To refresh the view, de-select, then re-select the Hide Empty Rows icon.

See also: ["Inspecting and Editing" on page 68](#).



Note: The overall set of event-attribute fields is defined in [Data Fields](#), but you can make or use custom subsets with the Field Set Editor (see ["Field Sets" on page 999](#)). Choose a set name to see only that predefined set of fields.

Events

Events begin at network devices that can sense and record instances of security-sensitive activity. Examples include a database record change, a syslog entry, a firewall transit, a router access, or scanning a door access card.

Such initial events are typically recorded in logs, and are sometimes called **base** or **raw** events.

When numerous source devices are reporting large volumes of relatively similar events, it is desirable to funnel these events through central event **concentrators** that forward a much-reduced set of representative or summary events.

When these events reach ArcSight [SmartConnectors](#), several things can happen.

- All received events are **normalized** (restructured) to make their information consistent and ready for analysis.
- All received events are **categorized** (appended with classification information) using ArcSight's event categorization taxonomy.
- If appropriate and the SmartConnector is configured to do so, events are **aggregated** to issue fewer and more meaningful events and to reduce network traffic.
- If appropriate and the SmartConnector is configured to do so, selected events are **filtered** out, to eliminate them as a further traffic or processing burden.
- For certain devices, the option may be available for the SmartConnector to apply analysis rules to incoming events and to issue **correlation** events concerning them.

At SmartConnectors, filtering **removes** events from the system. Aggregation **replaces** events with fewer new ones bearing summary information.

When the events from SmartConnectors pass to Managers they can again be considered **base** events in the sense that they are in a state prior to processing. More specifically, any event that is subject to further processing, even if the result of previous processing, can be considered a base event.

All base events entering the Manager are subject to:

- **Correlation** to derive more intelligence from the events. Correlation **adds** new events containing the results of correlation activity. You apply correlation through the rules and data monitors in their respective resource trees of the Navigator panel. Correlation events have flash icons in grid views.
- **Filtering** to selectively see and report on events. Filtering within the Manager does not actually discard events. You apply filtering with the resources in the Filters tree in the Navigator panel.

Note that all aggregation actually occurs at SmartConnectors, not within the Manager. You apply aggregation through the resources in the [Rules](#) tree of the Navigator panel.

There are only **base**, **aggregation**, and **correlation** events. It is important to note that any such event in the system can (if the right rules and data monitors are present) become the input to produce new correlation events. You should also note that the Manager's rules engine is designed to prevent infinite loops.

Apart from the events that originate on the network, and the correlation events the Manager issues in response to them, the Manager generates many other events of its own for a variety of purposes.

These **internal events** can be divided into [Audit Events](#) and [Status Monitor Events](#). You can use audit events to track, or react to, system **activity** at all levels of operation from data monitors to the database. Status monitor events are valuable for getting system **state** information. Review these topics on Audit Events and Status Monitor Events to become familiar with the characteristics of all the available events.

You can apply all analytic tools to any events present, whether base or correlation, originating externally or internally.

Event Annotation Fields

Event Annotation Fields

Event Annotation Field	Usage
Stage	Click this field to choose a different disposition state for the events' collaboration cycle. The default stages run from Initial to Closed. If you created your own stages as described in "Creating or Editing Stages" on page 300, these custom stages would be displayed here.
Assign to	Click this field to choose an ArcSight user to take the next step.
Is Reviewed	This read-only field tells you whether this event has been reviewed.
In Case	This read-only field tells you whether these events are already part of an ArcSight case. If they are, you have more ways to track their disposition. See "Working with Events in Cases" on page 640 for related information.
Correlated	This read-only field tells you whether these events are part of a correlated event chain. If so, you can learn more through the rules authored to control that chain of correlation. Note: You can configure the ArcSight Forwarding Connector to send correlation events along with the correlated base events from a source Manager to a destination Manager. However, the forwarded base events display the Correlated flag only in memory, which you can view on an active channel. If you want the forwarded base event's Correlated flag to persist in the database, set this property in <code>server.properties</code> on the source Manager: <pre>logger.base-event-annotation.enabled = true</pre> For instructions on how to edit the <code>server.properties</code> file, refer to the <i>ESM Administrator's Guide's</i> topic on "Managing and Changing Properties File Settings."
Hidden	This read-only field tells you whether these events are hidden from all but the assigned users of this stage.
Closed	This read-only field tells you whether the investigation of these events has been marked as closed. Closed events may no longer be visible to interested parties through active channels, for example.

Event Categorization

Events from unsupported or custom devices can generate events that the provided connectors do not know how to categorize. For example, if your organization has developed and deployed ArcSight FlexConnectors to collect and process events specific to customized network nodes, these custom events are not categorized per the usual method.

From the ArcSight Console, you can manually apply categorization to one or more custom events from a FlexConnector (or other custom or unsupported device). Once you apply categorization to events from a particular device (and its associated connector), the categorization is automatically applied to other events of the same type.



Note: To categorize events for an active channel using the task bar, you must add the **Device Event Class ID**, **Device Vendor**, and **Device Product** columns.

To apply event categorization to one or more events:

1. Select one or more of the same type of events that you want to categorize.
2. Select one or more events and choose the **Categorize Event** command from the System menu (or click the  toolbar button).
3. Select values from the given categories from the drop-down menus.
4. Click **OK** to apply the categorization information to events of this type.

This generates a SmartConnector update file (.csv) containing the new categorization files on the Manager. The Manager polls for new SmartConnector update files every 5 minutes, and updates the SmartConnectors when it finds new .csv files. So, within 5 to 20 minutes after you apply event categorization, new events of the same type are categorized in the same way.

Note that if a certain type of event is already categorized, this custom categorization has no effect. Otherwise, the custom categorizations take effect on all events of the same type going forward.

If you want to also apply the changes to previously categorized events, copy the newly created update file (.csv) to the SmartConnector installation folder as follows:

Copy from:

```
$ESM_HOME/user/agent/acp/categorizer/current/<vendor>/<product.csv>
```

Where \$ESM_HOME is the installation folder for your ESM environment.

To:

```
$AGENT_HOME/user/agent/acp/categorizer/current/<vendor>/<product.csv>
```

Where \$AGENT_HOME is the installation folder for your SmartConnector.

Event Handling Stages

Events coming into ESM go through the following stages:

1. Pre-persistence stage
2. Persistence stage
3. Post-persistence stage

Pre-Persistence Stage

At this stage, raw events arrive from SmartConnectors. ESM evaluates these events against the Priority Formula and enriches the events with Asset and Network Model information. Pre-persistence rules look for matching events at this stage.

See ["Priority Calculations and Ratings" on page 1019](#) for information about the priority formula, ["Assets" on page 830](#) for information about asset and network models, and ["Rule Types" on page 508](#) for information about pre-persistence rules.

Persistence Stage

The events are kept in active memory and accessed by live active channels. The filter associated with an active channel is applied to narrow down the events to be displayed on that channel. Event-matching for rules occur at this stage, and when events meet rule conditions, correlation events are generated. Then the prioritized, enriched events and correlation events are persisted in the database.

Post Persistence Stage

The rule engine triggers rule actions at the beginning of this stage. The events do not change at this stage, and are therefore removed from active memory since their copies are saved in the database. Queries, query viewers, and reports get their data from the database.

Field Sets

Field sets are named subsets chosen from the available [Data Fields](#). Field sets can help you quickly focus a grid view, Event Inspector, or other field array on a particular context such as customer accounts or vulnerability.

Field sets are a shareable resource that you can manage and apply through the Field Sets resource tree in the Active Channels section of the Navigator panel. (In the Navigator, choose **Active Channels**, and click the **Field Sets** tab.) These field sets also support the [Variables](#) data fields. Field sets supercede and include the previous concept of column sets.

There is a default list of field sets for out-of-the-box use, and to serve as examples.

See ["Creating a Field Set" on page 601](#) for information on how to create custom field sets, modify existing ones, and share them with other Manager administrators or operators.

See ["Sortable Field Sets" on page 1054](#) for information on creating and using sortable field sets.

See ["Using Field Sets" on page 899](#) for information on how to access field sets to build conditions.

Filters

Use filters to specify criteria that narrows the scope of monitored data and reduces the number, or constrains the nature, of the events displayed through the ArcSight Console. Filtering criteria are based on the Console's event [Data Fields](#), used in various combinations and with various conditions placed on their content. As you apply more restrictive filter parameters, the number of events reaching the Console may decrease, but the likelihood increases that the events are significant.

For example, you can create a filter that contains every firewall for the western region of the United States, and create another filter that contains every Intrusion Detection System (IDS) for the same region. You can also be more specific, by creating a filter wherein you only want to view firewalls and IDSs with certain IP addresses because they are labeled as **suspicious** IP's or IP's that may pose a possible threat to an enterprise. On the other hand, you can create filters that only contain networks that are labeled as **friendly** and seem to pose no threat at all, but you still want to monitor them. For display purposes, you can select a unique color for any filter. If an event matching the filter's conditions is generated, the event appears in the grid view in the specified color.

Applying filters to get optimum results is a core skill for network security analysis. While it isn't possible to anticipate specific solutions here, you should know the most efficient way to use the ArcSight Console's filtering tools.

Filtering Options

In the ArcSight Console, filtering is available in multiple ways, and how you choose to use these options can have a significant effect on your ability to precisely, flexibly, and rapidly author new analyses over the long term.

The primary event-filtering options are:

- **Filters resources:** The Navigator panel's Filters resource tree is (or should be) your master repository for filtering solutions. Using the Filters tree is the best way to work out an organized filter library. You can and should use the filters you develop here, through the Filters Editor, in other resources such as active channel views, reports, or rules. You can even use filter resources in other filter resources. By basing your solutions on hierarchical, resource-based filters, you gain the type of leverage granted by style sheets.

- **Active Channels resources:** The active channel resources in the Navigator panel can each store an individual filtering solution that is unique to a given channel or based on a Filters resource. When you use an existing active channel to create another, you carry forward and perhaps modify its filter.
- **Active Channel Editor:** You use the Active Channel Editor to create or modify the filters in individual active channels. Changes you make to active channels through this editor are limited to those channels and channels created from them. Such changes shouldn't be considered long-term or enterprise-wide.
- **Data Monitors:** Data monitors include a Restrict by Filter attribute that enables you to select a filter and therefore focus on the events to be displayed. See ["Data Monitors" on page 959](#) for more information.
- **Scheduled rules:** Rule schedules are configured for a group of standard rules. Refer to ["Scheduling a Rule Group" on page 559](#) for more details.
- **Inline view filters:** In any active channel grid view you can use the fields of the grid's top line to select filtering event-attribute values for certain columns, which will be used with implied AND operators to impose *ad hoc* filters. These filters are not retained with the prior active channel, but you can give the revised channel a name and save it through the Active Channel Editor.
- **Event-based filters:** Another quickly applied and contextual category of event filtering is offered by the event-attribute **Analyze in Channel** command. When you right-click an event attribute in a grid view you can choose Analyze in Channel and one of several filtering options that vary based on the data involved. Like inline filters, **Analyze in Channel** filters apply only to the current view and are temporary unless saved in a different named view.
- **Inline Filters:** You can add an inline filter to a channel view by clicking the Edit Inline Filter () button at the top right of the grid view to display the inline filtering fields. For more information see, ["Filtering Active Channels with Inline Filters" on page 234](#).
- **Threat Detector:** Threat Detector is a separately-licensed feature that leverages filter resources to configure profiles of interest. Profiles include a Restrict by Filter attribute that enables you to select the appropriate filter for the profile. See ["Threat Detector" on page 1013](#) for more information.

You should always remember that your most primary filter is the one imposed by your system administrator. Each user operates under the constraints of the access control lists (ACLs) configured for their user identity. These ACLs automatically filter out some portion of the total available event flow before it reaches you. Any filter you use or create adds to this fundamental constraint.

For more about putting filters to work, see ["Filtering Events" on page 306](#).

Global Variables

You can create variables that derive unique values from existing data fields, which you can apply locally in the resource you're working on to make monitoring and correlation more specific to particular scenarios.

In addition to these local variables, there is a global variable resource, which makes it possible to define a variable once, then re-use it in multiple places wherever conditions can be expressed (active channels, rules, filters, data monitors, and queries), and wherever fields can be selected (CCE, field sets).

Global variables are centralized and reusable, which make them an essential building block for user correlation in the Actors feature, and other advanced correlation scenarios.

Once created, global variables appear in the [Common Conditions Editor \(CCE\)](#) as additional fields on the Filters or Conditions tabs, as Group By arguments for data monitors and queries, and in rule conditions and actions. You can add variables to field sets in the Field Set Editor to extend the event and resource schema with values derived from other data fields.

The global variables feature also makes it possible to easily promote local variables defined for a particular resource into a global variable, where it can be re-used in other condition statements.

For details about using the Global Variables feature, see [Global Variables](#).

Grid View

A grid view is a type of view in the ArcSight Console that shows events summary information organized in rows and columns, or other types of information such as for certain [Resources](#). As new events occur, they are inserted at the top of the grid as new rows. Rows contain events while columns contain data fields. You can learn about working with grids in ["Using Views" on page 231](#).

IP Address Ranges

Use the following guidelines to input IP address ranges in a single string.



Caution:

- If you are defining conditions for IP address ranges in filters, queries, and rules, do not mix IPv4 and IPv6 addresses within the same IP address range.
- If you are defining your asset model and network zones, you may have mixed-family IP address ranges. However, OpenText does not recommend this. You might find the results confusing. See ["Managing Zones" on page 168](#) for related information.

Two-address range	A two-address range is in the format <code>firstAddress - lastAddress</code>, meaning any address between an arbitrary range of any two addresses, inclusive. IPv4 range: <code>192.168.0.0 - 192.168.255.255</code> IPv6 range: <code>2001:db8:fd0c:: - 2001:db8:fd0c:ffff:ffff:ffff:ffff:ffff</code>
CIDR notation (see RFC 4632)	The CIDR notation is in the format <code>address/prefix-length</code>. This format is more restrictive than the two-address range format where the range starts and ends. IPv4 range: <code>192.168.0.0/24</code> IPv6 range: <code>2001:db8:fd0c::/64</code>
Wildcard expressions	Fields on the right end of an address may be replaced with an asterisk, with no numeric data to the right of the first asterisk. The wildcard represents the range of all values for the field, from all-zero bits to all-one bits. This format is more restrictive than the two-address range format in where the range starts and ends. IPv4 range: <code>192.168.*.*</code> IPv6 range: <code>2001:db8:fd0c:*.*:*.*:*</code>

Inspect/Edit Panel

Located on the right side of the [ArcSight Console](#), the Inspect/Edit panel contains all the various [Resources](#) editors you use to create and modify analytic tools, as well as the Event Inspector you use to examine the contents of events. Using the Event Inspector and the resource editors is explained in the topics that relate to events and those resources.

See ["Events" on page 995](#) and ["Event Categorization" on page 997](#) for additional information about events.

Job Scheduler

You can schedule some tasks to occur automatically. Specifically, this feature is available for archiving and scheduling reports, Threat Detector snapshots, and rules. See these topics for information specific to scheduling jobs for particular resources:

- ["Scheduling Report Tasks" on page 461](#) for information on how to schedule reports individually or by group
- ["Scheduling Rules" on page 558](#)
- ["Trend Schedule" on page 475](#)
- ["Scheduling a Snapshot " on page 763](#)

This topic provides general information on how to schedule a job for any resource and view all scheduled jobs.


Where: Navigator > Resources > any resource group

To schedule a job:

1. Right-click a resource group you want to schedule and select **Edit Group**.
2. Click the **Jobs** tab in the Editor for a group.
3. Click **Add** on the Jobs tab.

This opens the Job Frequency dialog.

4. Define the schedule frequency (Hourly or Daily) and range of the job (Start and End dates, or indefinite). If your schedule frequency is Daily, additionally specify the time of day when the job should run.

The image shows a 'Job Frequency' dialog box with a blue title bar and a red close button. It contains two main sections: 'Schedule Frequency' and 'Schedule Range'. In the 'Schedule Frequency' section, there are three radio buttons: 'Hourly' (selected), 'Daily', and 'Weekly'. To the right of these is a field 'Every 1 hour(s) at 0 minutes after'. In the 'Schedule Range' section, there is a 'Start' field with a date and time '20 Jun 2016 18:35:35 PDT' and a calendar icon. Below it, there are two radio buttons for 'End': 'No End date' (selected) and 'End on' with a date and time '27 Feb 2017 14:15:05 PST' and a calendar icon. At the bottom are three buttons: 'OK', 'Cancel', and 'Help'.

5. Click **OK** to save the task and close the dialog.

Viewing all scheduled jobs

To view scheduled jobs:

1. Click the **Open scheduled jobs list** tool button.

The scheduled tasks are listed in the Viewer panel under **Current Jobs**.

2. Click a job in the list.

The status of previous and pending runs for that job are shown in the Scheduled Runs for *<Taskname>* list on the bottom part of the Viewer panel.

If a user adds scheduled resources to the Public folder, the scheduled jobs for these resources will not be deleted even after that user is deleted. The tasks are visible in Current Jobs but without an owner. The tasks without owners never run. If you do not want to see these non-functioning scheduled jobs, delete the scheduled resources created by users who have been deleted.

If scheduled jobs for resources are under the creator's personal folder, then these are automatically deleted when that user is deleted.

Troubleshooting Tips

If the Manager system clock time is changed after jobs are scheduled, some scheduled jobs might be kept in pending status temporarily, and/or synch back up on subsequent scheduled run times. It is best not to make dramatic adjustments to the system clock on a Manager on which critical jobs are already scheduled. If this problem occurs and does not correct itself soon enough, stop the Manager and re-set the system clock to the original date/time, or re-schedule the jobs.

Knowledge Base

The Knowledge Base is a problem-solving database that can contain information on event data, associated if-then-else rules, cases, and so forth. All information is derived from community expertise within your enterprise or based on your internal practices and policies.

Compare Knowledge Base articles to [Reference Pages](#), which provide built-in reference information about certain resources.

When you create a Knowledge Base article, you provide a URL or directory path to a specific vulnerability or exposure. You can add notes to Knowledge Base articles to relay information about the article. Using a note, you can write reminders, messages to the next shift, or any related information. Articles are displayed in the ArcSight Console, with associated links and article information. Knowledge Base articles are stored in these default groups:

- **Shared:** lists Knowledge Base groups and articles to which the logged-in user has access.
 - **All Knowledge Base:** lists all user Knowledge Base groups and articles.
 - **Personal Knowledge Base:** lists each user's own Knowledge Base groups and articles.

- **Public Knowledge Base:** lists Knowledge Base groups and articles accessible to all users.
- **Unassigned:** lists Knowledge Base articles that do not belong to a group.

Logical Operators

This topic describes the logical operators you can use in condition statements. Certain operators do not appear in circumstances where they are not applicable.

Entering IP address ranges

The insubnet operator uses a range of IP addresses. Use the following guidelines to input IP address ranges in a single string.



Caution: If you are defining conditions for IP address ranges in filters, queries, and rules, do not mix IPv4 and IPv6 addresses within the same IP address range.

Two-address range	A two-address range is in the format <code>firstAddress - lastAddress</code>, meaning any address between an arbitrary range of any two addresses, inclusive. IPv4 range: <code>192.168.0.0 - 192.168.255.255</code> IPv6 range: <code>2001:db8:fd0c:: - 2001:db8:fd0c:ffff:ffff:ffff:ffff:ffff</code>
CIDR notation (see RFC 4632)	The CIDR notation is in the format <code>address/prefix-length</code>. This format is more restrictive than the two-address range format where the range starts and ends. IPv4 range: <code>192.168.0.0/24</code> IPv6 range: <code>2001:db8:fd0c::/64</code>
Wildcard expressions	Fields on the right end of an address may be replaced with an asterisk, with no numeric data to the right of the first asterisk. The wildcard represents the range of all values for the field, from all-zero bits to all-one bits. This format is more restrictive than the two-address range format in where the range starts and ends. IPv4 range: <code>192.168.*.*</code> IPv6 range: <code>2001:db8:fd0c:*:*:*:*</code>

Logical Operators for Conditional Statements

Logical Operator	Description
=	Equals Use this operator when the entire string is known, such as for an event Name or User name.
!=	Not equals Use this operator to exclude one or more known values, such as events involving a specific network domain or user.
<	Less than
<=	Less than or equal to
>=	Greater than or equal to
>	Greater than
Between	Between a range of comma-separated, comparable field types; for example, numeric types, date, date and time, IP address, or MAC address types. The minimum value for the range must appear first in the expression before the maximum. For example, to match a value of 20, use between (-100,100) If your field type is IP address, the IP addresses must be of the same family: IPv4 or IPv6 addresses.
BitAnd	Equals, for bitmap fields
Contains	Contains the specified substring Use this operator to exclude a large set of events, such as all events whose name contains "virus." Use this operator with caution as it is relatively slow to evaluate and prone to matching more events than you intended.
ContainsBits	Returns true or false. The ContainsBits operator applies to event fields that are bit vector data types, that is, fields that combine a set of independent Boolean flags. The right side value is a list of applicable named flags. The operator evaluates to True if the value of the selected flag is set for the event. For example, <code>Event Annotation Flags ContainsBits "correlated;inCase"</code> returns true if both of the annotation flags Correlated and inCase are set for the event being investigated.
ContainsList	Compares one list to another to see if the second list is a subset of the first. For example, (a, b, c) ContainsList (a, b).

Logical Operators for Conditional Statements, continued

Logical Operator	Description
ContainsValue	Returns true or false Use this operator where the left-hand-side operand is a list of some data type, and the right-hand side operand is a single value (field or literal) of the same data type. For example, a variable <code>IPList</code> is a list of IP addresses obtained from a multi-mapped active list, with values {192.0.2.0, 192.0.2.1, 192.0.2.2}. You then use <code>ContainsValue</code> in the following conditional statements: <pre>IPList ContainsValue TargetAddress</pre> This condition returns true or false depending on whether <code>TargetAddress</code> value is contained in <code>IPList</code>. <pre>IPList ContainsValue 192.0.2.0</pre> This condition returns true because the right-hand side value is contained in <code>IPList</code>. <pre>IPList ContainsValue 192.0.2.24</pre> This condition returns false because the right-hand side value is not contained in <code>IPList</code>. You can use <code>ContainsValue</code> in both queries and in-memory resources.
EndsWith	Ends with specified substring. Use this operator for domain names. For example, you might want to match events involving the .mil domain.
EqualsList	Compares one list to another (for example, active list, session list). If the two lists have the same entries, the statement evaluates to true. For example, <code>(a, b, c) EqualsList (a, b, c)</code>.
In	Standard SQL operator for membership test You can use a comma-separated list for the right operand. For IP addresses, the left and right operands can be mixed-family addresses.
InActiveList	Event appears in the specified active list. <code>InActiveList</code> operates on items in the event and actor schemas. It does not evaluate items in other non-event schemas (such as cases or assets). For example: <pre>(a, b, c) InActiveList (a, b, c, d)</pre> Note: The <code>InActiveList</code> operator only evaluates single-value attributes, and treats multi-value attributes, such as Actor Account ID and Role, as single-value attributes.
InGroup	Tests for membership in a specified category.
InList	Determines whether a given item is in a list and, if so, evaluates to true.
InSubnet	For IP addresses in the specified subnet Caution: The IP addresses must be in the same family, for example, pure IPv4 or pure IPv6 addresses only. IPv4-compatible and IPv4-mapped IPv6 addresses are processed as IPv4 addresses. See examples of IP address ranges here.

Logical Operators for Conditional Statements, continued

Logical Operator	Description
Is	Tests true for the selected state, null or not-null Use this operator to test whether or not a value has been supplied. You would use this in rules to tell the difference between a string that does not match versus a string that was not supplied. For example, you could use this to find all events that were missing their event names.
IntersectsList	Compares one list to another. If the two lists have one or more entry in common, the statement evaluates to true. For example: (a, b, c) IntersectsList (b)
Like	Standard SQL operator for simple pattern matching for string type: "_" wildcard for single character; "%" wildcard for multiple characters
Matches	For extended regular expression pattern-matching for string types using Perl 5 syntax Supports regular expressions (regex). Note that Matches is used in rules only.
On	Event occurs on this date
StartsWith	Starts with specified substring Use this operator for testing URIs such as event categories or resource locations (for example, Customer or Connector locations in their respective Navigator trees), or to test the root of a hostname (for example, if your web servers are named WebServer1, WebServer2, and so forth, you could use "hostname startsWith WebServer").

Managed Security Service Providers (MSSPs)

Managed Security Service Providers (MSSP) can use slicing and dicing query-trend approaches to create focused reports for multiple customers built from what are initially broad range queries.

Manager

The Manager is the component that manages, cross-correlates, filters, and processes all security-event occurrences in your enterprise. The Manager includes CORR-Engine for storage management, a Correlation Engine, Connector Data Manager, tracking and resolution functions, and analytics and reporting capabilities.

Navigator Panel

Located on the left side of the ArcSight Console, the Navigator panel contains all the trees you use to organize analytic and operational [Resources](#), tools, and targets. These resources come in many types, such as active lists, connectors, rules, and users, all of which are summarized in the topic ["Navigating" on page 43](#).

Notifications

"Notifications" usually refers to the event-related messages sent to e-mail addresses or cell phones.

Sending notifications is one among several rule actions that can be performed when a rule is triggered (See ["Rules" on page 1034](#)). When you create a rule and add a Send To Notifier action, you will be able to select the notification group that will receive the message. For more information on rule actions, see ["Managing Rule Actions" on page 532](#).

The key entities in the notification structure are Notification Groups, Escalation Levels, and Destinations.

Notification Operation

When a rule that has a notification action triggers, the notification engine notifies all **active** destinations in the first escalation level within that group. The notification engine then waits for a certain time period to receive an acknowledgment to that notification.

You can acknowledge notifications by any one of these methods:

- Reply to the e-mail or cell phone message
- Click the **Notifications** button () in the Console's toolbar to use the Notifications Manager in the Viewer panel. (See ["Managing Received Notifications" on page 216](#).)

The length of time that the notification engine waits for acknowledgment depends on the event severity, and can be configured through the context (right-click) menu's Wait Time setting.

If no acknowledgment is received within the specified time interval, the same notification is escalated to the next level within the group.

This process repeats until there are no more escalation levels or the notification is acknowledged by any of the recipients. The one exception to this procedure is the escalation

procedure carried out for **informative** notifications (the Informative option was set while defining the notification action in the rules editor). In this case, notifications are only sent to the first escalation level in the group and do not require acknowledgment.

SMTP is used to send e-mail. An SMTP server must be configured either at install time or through Context (right-click) menu e-mail settings. For notifications, the relevant fields are "from address", which designates the e-mail address of notification e-mail sent from the system, and the "outgoing e-mail server," which is the SMTP server used to send e-mail. It is important to ensure that the "from address" specified is one that will not be rejected by the SMTP server, since some SMTP servers will reject unknown e-mail addresses.

POP3 and IMAP can be used to check for e-mail acknowledgments. You can specify these options at install time, or through Context (right-click) menu e-mail option settings. For acknowledgements, the relevant fields are "incoming mail server," which is the POP/IMAP server to specify to check e-mail, "incoming mail protocol," which is either POP3 or IMAP, "account" and "password," which are the login name and password to access the mailbox from the incoming mail server. Note that replying to mails from the notification "from address" should reach the mailbox accessible to the "account" login.

Notification Groups are the point of interface between the rules engine that specifies the notification action and the notification engine that sends out the notification. Within each notification group, there can be any number of escalation levels. Each escalation level can contain multiple destinations.

The following groups can help you manage and organize groups and destinations.

Notification Groups and Destinations

Destination Group	Purpose
Shared	Notification groups and destinations to which logged-in users have permission.
All Destinations	All groups and notification destinations (only Administrators have permissions to this group). Administrators who have inspect and edit permission on the All Destinations group also have permission to change notification settings.

Testing Notification Escalations

Escalation procedures are tested by generating an internal Low Severity event. This event triggers the escalation within the group tested as though a real Low Severity event occurred. Notifications are immediately sent to all destinations within the 1st level (1). If 1st level destinations do not respond to the notifications within the set wait time for Low Severity events (default is 2 hours), the test notification escalates to the 2nd level (2), and so on.

Notification Destinations

Notifications are sent to destinations. Notification destinations may optionally be associated with a user, and when that is done, destination information, such as e-mail address or cell phone, is automatically populated from the user's profile. You can also change the user's destination information without changing the user's profile.

Each destination can be an e-mail or cell phone contact and have an associated start and end time, which is the time period during the day when the destination is expected to be active. Each destination can also be optionally associated with a user. When the destination receives a notification and the user is logged into an ArcSight Console, the user is notified through the notification status button on their display.

Notification destinations can be managed with drag and drop functionality. You can move or copy notification destinations into escalation levels within the same or other notification groups from the Administration window. If a group is deleted, the destinations within that group are also deleted.

 **Note:** To copy multiple resources at once, use Copy and Paste. You can drag and drop only one resource at a time.

Notification Acknowledgements

Once you receive a notification, it is important that you acknowledge it within the allotted time window, to prevent automatic escalation to the next-level notification destination.

Immediately acknowledging and resolving significant events is crucial to securing any enterprise. Use the Console's Notifications Manager to check status and help resolve issues.

See also ["Managing Received Notifications" on page 216](#).

Packages

A **Package** is a resource that contains a set of related resources. A package of resources can be installed or unloaded as a unit. ArcSight Solutions are delivered as packages, but you can create your own packages, as well.

A **Bundle** is a file (with extension .arb) that contains one or more packages. You can import and export bundles and install and uninstall the packages that the bundles contain. When you import a bundle, the source file is saved as a file resource (see ["Managing File Resources" on page 727](#)). You can view the original package contents (the package archive) or the current package contents at any time.

An **uninstalled package** is a package that has been imported or created, but not yet installed in the system resource tree (see ["Resources" on page 1032](#)). Packages that have been installed can also be manually uninstalled. The default behavior is to install the package when it is imported.

When a package is deleted, the resources it contains can be left in the system resource tree or they can be deleted along with their package.

Packages can have dependencies on other packages or on features such as Threat Detector. Two ArcSight Solution packages may share a third package in common, for example.

See also ["Managing Packages" on page 730](#).

Threat Detector

ArcSight's Threat Detector can detect subtle, specialized, or long-term patterns that might otherwise go undiscovered in the flow of events. This topic discusses pattern concepts. See ["Using Threat Detector" on page 747](#).

A pattern is a distinct, repeating network transaction (event) that is uniquely identified by its source and target IP addresses. Patterns are further qualified by the involvement of selected attributes such as event names or categories. There are, of course, many such patterns and most are normal or benign. The point is to establish and mask out normal traffic in order to let new or atypical traffic stand out. Separating "signal from noise" in this way makes possible very early (day zero) detection and very subtle (low and slow) detection. Once detected, such traffic can be analyzed or responded to.

Threat Detector uses a **profile** to specify potentially qualifying events on the basis of attributes and time spans. When you apply a profile, manually or on a schedule, it captures a **snapshot** of the events that did qualify, on the basis of raw associations. The contents of snapshots are then reviewed by an analyst to identify event **patterns** to explore in pattern views or the Pattern Inspector.

You define profiles in the Profile Editor in the Inspect/Edit panel. You manage your profiles, snapshots, and discovered patterns through the Profiles, Snapshots, and Patterns tabs of the Navigator panel's Patterns resource tree.

You use the Viewer panel to observe the graphical results of executed snapshots and the patterns those snapshots discover.

Pattern Concepts

A pattern can be any recurring relationship between one or more pairs of source and target IP addresses, that you deem to be significant in relation to certain event attributes. You can regard the patterns you discover as benign or hostile, depending on your policies and postures.

Event-pattern profiles are also constrained by Start and End time limits, filters, and by minimum numbers of associated events (pattern length) and times discovered (occurrences, or pattern support).

Once captured in snapshots, you can examine the event data as raw association information in graphical snapshots, or as graphical patterns in the Patterns tab of the Patterns resource tree.

Each box in a pattern view represents one pattern. The line items in the box are the individual events that were discovered to have associations. Each event component of a pattern (box) relates to the chain of links from which the pattern was derived, in the visual snapshot.

A snapshot view is a graphic hierarchy of related event nodes. The "support" value for each node is the number of times that event occurred in conjunction with its related events. This overall hierarchy is a raw presentation of the events, useful for analysts but not meaningful to operators.

The discovered events all share the attributes specified in the profile. The pattern-discovery process first tests for equality in the values found for the specified attributes. Secondly, it tests for a selected transaction scheme. When the specified minimum number of event relationships reoccur, a pattern exists.

Discovering Patterns

Patterns are identified by first dividing the event stream into multiple transactions. For example, all of the events with a given source and target IP address may constitute a transaction - they represent all the traffic flowing from that source to that target. It may also be helpful to cluster transactions into super-transactions to identify patterns that involve cascading exploits toward multiple devices (that is, device A attacks device B which, in turn, attacks device C).

The events occurring in each transaction are then characterized using a subset of the event fields (for example, the event name or the event category).

Finally, events that frequently occur together in multiple transactions are identified and grouped together. These events are further sub-grouped by support level. For instance, events A and B may occur together 2,000 times while events C and D occur in the same transactions but only 10 times. Threat Detector would create two patterns in this case: one for A and B and

a second one for C and D. To give another example, events F, G, and H may occur together in the same transactions 100 times while F and G occur without H in 5 additional transactions. All of these occurrences would be rolled together into the same pattern. F and G would have a support of 105 while H would have a support of 100.

Pattern Analysis

Pattern analysis, overall, falls into two basic phases: initial collection, identification, and sorting, and on-going routine processing.

Initial Phase

To accomplish phase one, you generally use broader profiles and more frequent snapshots in an attempt to capture examples of **all** the patterns that appear in your networks.

Once collected, there is a period of initial analysis in which you identify the patterns that are normal or benign. Making these evaluations requires in-depth knowledge and familiarity with the traffic in your enterprise, as well as using the analysis tools. There is no set procedure for this basic collecting and sorting process.

However, the best method for moving officially "uninteresting" patterns **out** of the analysis workflow is to use annotation. While it is possible to use filters for this purpose, it is more reliable to move patterns by annotation to a stage such as **Closed** because this assures that the pattern has actually been inspected and classified.

Routine Pattern Processing

In an environment where the routine event patterns are mostly known and appropriately classified, you focus on the new and as-yet unclassified.

The basic approach to routine pattern analysis consists of two phases: managerial (or triage or workflow initiation), and analysis.

Workflow Management

As Threat Detector turns up new or unclassified patterns, a designated user needs to review them and start them through the workflow.

Newly discovered patterns are handled by using the Annotations feature to assign them to a stage such as **Follow-up**, or simply **Closed**, and optionally to a particular user.

Specific procedures and decisions, of course, depend on the internal processes of your enterprise and the patterns encountered.

Pattern Analysis

As an analyst dealing with day-to-day pattern discoveries, your basic process can be as follows.

Using the appropriate filters, view the patterns that are new and assigned to you in the Pattern Inspector.

Review these patterns in the Pattern Inspector and compare their transactions historically to those found in other snapshots, using the Snapshot menu.

Use the Show Related Events feature to gain more intelligence about the sources and targets that appear in the patterns.

Remember that events in a grid view are subject to all the ordering, graphing, filtering, reporting, and inspection tools available in the Viewer panel.

Visualize the source and target relationships using Show Event Graph.

Pattern Disposition

Acting on reviewed patterns can include:

- Assigning a new stage or user

The pattern may need further analysis or some other handling, by another user, or can simply be closed. Use the **Annotate Pattern** command to make this disposition.

- Creating a rule

If a pattern represents activity that needs to be reported, monitored, evaluated, or otherwise acted upon automatically, use the Create Rule command to build a rule based on the pattern's items.



Note: Remember to express an appropriate Time Frame value in the Aggregation tab of the Rules Editor. The scope of a rule's time frame is critical to its effectiveness.

- Deploying a rule

Once created, if a rule is of value to the enterprise, you should copy or move it to the Rules/Shared/All Rules/Real-time Rules group in the Navigator panel's Rules resource tree.

Threat Detector Expertise

On a work-a-day basis, the following points will help you make the best use of Threat Detector.

Workflow

Threat Detector analysis may also be scheduled. For example, once per hour the prior hour may be analyzed using three different profiles. The patterns discovered by each profile will be stored in a designated group in the Patterns resource tree.

Each pattern also has certain annotation features associated with it that will be familiar to users of trouble-ticket systems. Each pattern can be flagged as being at a given stage (for example, Queued, Acknowledged, Under Investigation, Under Observation, Normal Activity, and so forth). Patterns may also be assigned to an user for further investigation.

Initially, many new patterns will be observed and will need to be characterized. Does the pattern represent a threat or is it a result of normal activity on the network? Should a rule be generated? Or is more observation of the pattern required in order to understand it?

Over time, only a few new patterns will be observed each day. These will be delivered in the Queued stage. In the simplest workflow, the operator must resolve these patterns or assign these patterns to others for resolution each day.

When patterns are observed again, you can set it up to either quietly mark the pattern as observed again or to bring the pattern to the attention of the operator.

Visualization

Event graphs have a clustering ability that makes them very useful when illustrating the interactions represented by a pattern resource.

Suppose events F, G, and H occur together in the same transactions 100 times while F and G occur in 5 additional transactions. All these occurrences would be rolled together into the same pattern. The event graph would cluster the 100 sources where F, G, and H occur together. It would also cluster the sources where only F and G occurred.

To use a somewhat more concrete example, one cluster might represent a Nimda Worm's attempts to infect IIS installations. The second cluster might represent successful infections.

Applications

Threat Detector can be used to characterize the traffic on newly protected networks (for example, new customers for MSSPs, new divisions for large corporations, and so forth). It can also characterize traffic from new sensors.

Threat Detector is also a key element in the ongoing operation of an installation. Using periodic, scheduled analysis, operators can always be kept up to date as new event patterns appear. Frequently, these patterns will indicate new worm or exploit behavior.

Payload

"Payload" refers to the **information carried in the body of an event** network packet, as distinct from the packet's "header" data. (See ["Events" on page 995](#).) While security event detection and analysis usually centers on header data, packet payload (📄) may also be significant for historical analysis purposes.

Typically, devices discard payloads after a certain period of time. As described in ["Working with Event Payloads" on page 301](#), you can retrieve, preserve, view, or discard payloads using the ArcSight Console. Since event payloads are relatively large, they are not stored by default. Instead, you can request payloads from devices, for selected events, through the ArcSight Console. If the payload is still held on the device, the SmartConnector retrieves it and sends it to the Console. (See ["SmartConnectors" on page 1050](#).)

Payloads are downloaded and stored only on demand. Whether an event has a payload to store is visible in event grids. Unless you specifically request to do so, only the event's "payload ID" (information required to retrieve the payload from the event source) is stored. Payload retention periods are controlled by the configuration of each source device.

A payload that has already been downloaded and stored in the database can either be manually selected and deleted, or removed based upon the event-retention policy.

If the payload's format is not recognized by the database, its data will not be lost; instead it appears "unparsed" in the event. The event name attribute generally contains the complete data in this case.

Prioritization Fields

Events include fields whose values help you evaluate each event's overall priority and importance, and determine which events you should investigate first. The prioritization field values take into account a number of factors including:

- Vulnerability of the Target Asset
- Active List Contents
- Open Ports on the Target Asset
- Asset Criticality

Event Prioritization Fields

Data Field	Description
Model Confidence	Is the target asset modeled and if so, to what degree? This factor depicts the confidence we have in our model. This value depends heavily on whether target assets of interest are modeled in the system. If the only data point for an asset is its ID, then it is likely that this is either an asset range, or an asset that was modeled manually. The fact that the target asset is in the system at all provides some degree of model confidence. Model confidence is higher, though, if the target asset has been scanned for open ports and vulnerabilities.
Asset Criticality	How important is the Asset? This factor encompasses the criticality of the attacked asset.
Relevance	Does it appear probable that the attack succeeded? This factor performs an open port correlation (check to see if the target port is open) and vulnerability correlation (check to see if one of the exploited vulnerabilities is exposed).
Severity	How serious is this attack? This factor encompasses the severity of the event (ArcSight Severity), the severity of the exploited vulnerability (how much it is exposed), any user-supplied filter weighting, and the presence of the Source IP Address in various compromised and hostile active lists.
Priority	Should this event be investigated right away or not? This value is calculated by a formula that considers the values of the previous four fields, as described in the next topic.

Priority Calculations and Ratings

The priority formula, formerly referred to as the Threat Level Formula, is a series of five criteria that each event is evaluated against to determine its relative importance, or urgency, to your network. This topic describes the calculations used to determine an event's *priority rating* ("[Priority Rating](#)" on page 1024).

Priority evaluation is applied to all the events that the Manager receives from SmartConnectors. The event's priority lets security operations personnel know whether this is an event that warrants further notice. The priority value assigned to an event is essentially the severity the event was assigned by the original reporting SmartConnector, as modified by the weighting schemes model confidence, relevance, severity, and asset criticality. Each of these four criteria described in the table below contributes a numeric value to the priority formula.

Each of the four factors evaluates to a value in the range of 0 to 10, where 0 is low and 10 is high. The values have a specific positive or negative influence (weight) on the original SmartConnector severity value. See "[Prioritization Fields](#)" on the previous page for definitions of these factors.

The priority formula consists of 4 factors that combine to generate an overall priority rating. Each of the criteria described in [Factors Contributing to Priority Evaluations](#) contributes a

numeric value to the priority formula, which calculates the overall importance, or urgency, of an individual event.

All values fall in the range between 0 and 10. A high priority factor generally indicates an event with a higher risk factor. Not every high priority event is necessarily a threat, however. For example, if a critical e-mail server fails, the priority of the events reporting it may be very high, although it does not necessarily represent a threat to your network.

The following table describes the factors considered in the ESM priority evaluation. If you require help in changing the values, enter a case with Software Support. The maximum score for each factor is 10: if the value of qualifying conditions for that factor totals more than 10, the amount over 10 is not considered.



Note: You can view an event's priority information by right-clicking the event on the grid and selecting **Debug Event Priority**. The window displays information on how the priority score was determined for the selected event. The values described in the following table come from actual values stored for the events. The debugging information, however, is real time without history.

If you set the severity through a rule action, the debug event priority shows this value; however, the debug information does not cover this particular rule action. This is because the values described in the information are based on actual values stored for the event. Event conditions defined in the rule are based on live evaluation of the current state of the system.

Factors Contributing to Priority Evaluations

Priority factor		Description
Model Confidence		Model confidence refers to whether or not the target asset has been modeled in ESM and what information the modeling revealed. Maximum score = 10.
	+4	Target asset is modeled in ESM and its asset ID is present. If these are the only data points present for the asset, this is likely an asset range or a system that was modeled manually.
	+4	Target asset has been scanned for open ports.
	+4	Target asset has been scanned for vulnerabilities.
Relevance		Relevance of the event to the asset is based on whether the event contains ports or known vulnerabilities and whether they are exposed. If an asset does not expose the vulnerabilities or ports in the event, the event is not relevant to the asset. Maximum score = 10.

Factors Contributing to Priority Evaluations, continued

Priority factor	Description
+5	Ports <pre> graph TD A{Event contains port?} -- No --> B((+5)) A -- Yes --> C{Asset scanned for open ports?} C -- Yes --> D{Port open on asset?} C -- No --> E((+5)) D -- No --> F((0)) D -- Yes --> G((+5)) </pre>
+5	Vulnerabilities <pre> graph TD A{Is there a known vulnerability mapping on file at the Manager?} -- No --> B((+5)) A -- Yes --> C{Asset scanned for vulnerabilities?} C -- Yes --> D{Vulnerability exposed by asset?} C -- No --> E((+5)) D -- No --> F((0)) D -- Yes --> G((+5)) </pre>
Severity	Severity is a history function: Has the system been attacked or compromised before, or has the attacker scanned or attacked the network in the past? Different scores are assigned based on the attacker and target's presence in one of ESM's threat tracking active lists (/All Active Lists/ArcSight System/Threat Tracking), whose contents are updated automatically by ESM rules. Maximum score = 10.
+6	The asset appears as an attacker in the active list /ArcSight System/Threat Tracking/Infiltrators List.
+5	The asset appears as an attacker in the active list /ArcSight System/Threat Tracking/Hostile List.
+3	The asset appears as a target in the active list /ArcSight System/Threat Tracking/Compromised List.
+3	The asset appears as an attacker in the active list /ArcSight System/Threat Tracking/Suspicious List.
+1	Asset appears as an attacker in the active list /ArcSight System/Threat Tracking/Reconnaissance List.

Factors Contributing to Priority Evaluations, continued

Priority factor	Description
Asset Criticality	Asset criticality measures how important the target asset is, as set by you in the network modeling process by using the standard asset categories /System Asset Categories/Criticality/Very High, High, Medium, Low, and Very Low. For example, customer-facing systems or devices with access to confidential information would be classified with a High criticality level, whereas a staging or test system might be Low. Maximum score = 10.
+10	The asset is found by the filter /System Asset Categories/Criticality/Very High
+8	The asset is found by the filter /System Asset Categories/Criticality/High
+6	The asset is found by the filter /System Asset Categories/Criticality/Medium
+4	The asset is found by the filter /System Asset Categories/Criticality/Low
+2	The asset is found by the filter /System Asset Categories/Criticality/Very Low
+0	The asset is not categorized with any of the above categories.

You can use asset aging to reduce asset confidence level as the time since the last scan increases. For information on configuring that, see "Asset Aging" in the [ESM Administrator's Guide](#).

The priority calculation formulas are made up of basic elements organized by operators called Sum and Difference. These elements are based on simple condition expressions.

More information:

- ["Priority Elements" below](#)
- ["Priority Operators" on the next page](#)

Priority Elements

The basic formula elements each return a positive numeric value or zero. Individual element values can be configured by changing the Value attribute associated with the XML element for each condition.

Some of the elements are predicates that test a specific condition. If the condition for a specific element is satisfied, these elements return a positive value; otherwise, the element returns zero.

Predicate elements can also be negated using the Negated attribute. In that case, they return a specified value if the condition is not satisfied, and zero if the condition is satisfied.

Priority Elements

Prioritization Element	Description
HasOpenPort	Takes a non-zero value if the target asset has a particular port open.
HasVulnerability	Takes a non-zero value if the target asset is vulnerable to the attack captured by the alert under consideration.
HasVulnerabilityMapping	Takes a non-zero value if the signature of the context event has not been mapped to a vulnerability.
HasValue	Takes a non-zero value if the specified event attribute has a value.
InActiveList	Takes a non-zero value if the target address belongs to one of the active lists whose URI is provided in the formula.
Constant	Evaluates to a constant non-zero value. It does not rely on event-specific conditions or any other variable; it remains constant, as the name implies.

Priority Operators

There are two aggregation operators used in the priority calculation formula, **Sum** and **Difference**. The **Sum** operator adds the values of all of the elements that it contains. The **Difference** operator subtracts the sum of all of the values of the subsequent elements from the value of the first element it contains.

Both operators have two attributes, **maxValue** and **weight**.

MaxValue Attribute

MaxValue is used to clip the result after the operator aggregation is carried out. After aggregating, the result is also normalized, which is achieved by dividing the result with MaxValue. For example, if we have an element like

```
<SUM maxValue = 100>
```

and it has two child elements, each of which evaluate to 80, the pre-normalization value will still be 100 and not 160. After normalization, the final result for this example will be 1. Similarly, there is an implied lower limit or minimum value of zero on these elements.

Weight Attribute

The Weight attribute is used to scale the result after operator aggregation and normalization are carried out. So, as in the example previously described, if the aggregating element was:

```
<SUM maxValue = 100 weight = 7>
```


the result after normalization is 1, and after scaling, it becomes 7.

Each of the formulas have an implied max**Value** of 10 since each of the four fields in the alert take values in the range 0-10 (inclusive).

Priority Rating

The priority of an event is a calculated overall rating based on **agentSeverity** adjusted by Model Confidence, Relevance, Severity, and Criticality using a detailed formula. (See "[Priority Calculations and Ratings](#)" on page 1019.) All four factors are fields in the event schema, and can thus be used in correlation.

The priority rating is color coded and displayed in [Active Channels](#). You can sort events in the grid view according to priority. Priority is a good basis for deciding what to look at first in your event monitoring workflow, and priority is one of many useful criteria on which to build filters, rules, reports, and data monitors.



Priority Ratings in Active Channels: The Priority column in the default live channel view shows the overall priority rating for each event based on calculations from the other five priority criteria.

The score and color scale used in the priority display are as follows:

Priority Rating Color Indicators

Priority	Color	Description
0-2	Green 	Very low. This event is likely a routine function, such as routine file access or a successful authentication by an authorized user. An event that may have started out with a higher priority can become very low priority when it is proved to have failed.
3-4	Blue 	Low. This event is likely a common function, such as a setting change or a scheduled system scan.

Priority Rating Color Indicators, continued

Priority	Color	Description
5-6	Yellow 	Medium. This event is a potential concern, such as pre-attack scan activity, policy violations, and identified vulnerabilities. Medium priority events are often hostile attempts whose success or failure is not confirmed.
7-8	Orange 	High. This event is a concern, such as attack formations, potential breaches, or misuse, including traffic to a dark address space, incorrect registry values, or a SYNflood.
9-10	Red 	Very high. This event is a grave concern, such as verified breaches or a DHCP packet that does not contain enough data. Items with a very high priority should be investigated immediately.

Queries

A query is a resource that defines the parameters of the data you want to report on derived from a data source. Queries are used in [Reports](#) either directly or as the basis for Trends reporting ("[How Trends Work](#)" on page 467).

Queries can use as a data source the database of events, cases, notifications, modeled network objects (assets), trend data, active list, or session list. Reports then bind data results from queries and/or trends into a display format based on a report template.

See "[Building Reports](#)" on page 388 for an overview of all reporting tasks and tools, including how to build queries or trends and how to use a provided templates.

Queries and Trends

You can use the result of a query as the basis for one or more reports or trends. For a detailed description of how queries and trends can be used together, see "[Query-Trend Relationships in Reporting](#)" on page 468.

Building and Running Queries

You can access queries and associated editors in the Reports resource in the ArcSight Console.

See "[Building a Query](#)" on page 320 for information on how to navigate to and use the Query Editor to define query settings.

Query Viewers


Query Viewers are a type of resource for defining and running SQL queries on other resources, including trends, assets, cases, connectors, events, and so forth. Each query viewer contains an SQL query along with other logic for establishing and comparing baseline results, analyzing historical data to find patterns in network activity, and performing drill-down investigation on a particular aspect of the results.

You can use query viewers to run the same queries used for reports, and get results quickly. Then, if desired, you can generate a simple report directly from the query viewer results. Full-featured reporting (with queries, trends, and templates) is still offered for more robust reporting requirements (see ["Building Reports" on page 388](#)), but query viewers provide a shortcut to running those same SQL queries apart from reporting.

Query viewers provide high-level summaries to monitor system health, reveal trends, and allow for drill-down investigation of all types of resources. Query viewers can work with trend tables rather than event tables, and so can return results much faster than [Active Channels](#).

See ["Query Viewers" on page 340](#) for information about using and building query viewers.

Also, you can generate simple reports directly from query viewer results.



Query viewers provide:

- **A quick way to run SQL queries and trends apart from full-scale reporting.** If you want to run a pre-built SQL query and view results quickly, or build and test several iterations of a custom query, query viewers are an easy way to do it. (You can also generate a simple report directly from a query viewer.)
- **High-level summaries.** For example, using the aggregation provided by queries and trends allows summaries of “interesting things” over the last month, day, or hour.
- **Non-event-based summaries.** Queries can be used to analyze resources other than events (such as assets and cases).
- **Event-based summaries.** Queries can be used to analyze events, and eventually lead to active channels (with drill-down investigation).
- **Baselines.** Analysts can apply a baseline to the information resulting from a particular run of a query viewer. A baseline acts as a reference point against which to compare results of other runs of the same query and highlights the *deltas* (differences) to help identify areas that vary significantly from normal.
- **Drilldown.** Query viewers can provide drilldown investigation into the same or another query viewer for good performance on the next level of results as well. Ultimately, the drilldown can lead to an event channel, where the performance costs are the trade-off for the power of event-based analysis in an active channel. The query viewer author defines the appropriate drilldown paths and levels.
- **Performance.** Query viewers can use trend tables which are typically much smaller than event tables, and can be pre-built with summary views in mind. So, in most cases query viewers can return and display results faster than [Active Channels](#).
- **History.** When based on trends, query viewer result data can be kept for as long as desired and be independent of the event archival process.
- **Flexibility.** ArcSight provides both pre-built query viewers and a resource editor for adding custom query viewers to suit the needs and environment of your organization.
- **Presentation Options.** Query viewer results can be displayed as tables (with baselines, if desired), pie charts, and bar charts, and added to [Dashboards](#) for quick display and monitoring.

Reference Pages

Certain [Resources](#) among those you find in the trees of the [Navigator Panel](#), or events you see in the Viewer Panel ("[Views](#)" on page 1096), have pointers to additional reference information. To check for this information, you right-click an individual event, resource, or resource group and choose **Reference Pages**.

If there are pointers available, you see the Reference Pages dialog box. Select one or more items and click **View** to open them in Viewer tabs. If no content is available, click **OK** in the "none found" dialog box.

Some reference page pointers are pre-populated. You can edit these, or add new references, through the **Group Editor** (as described in ["Working with Resource Groups" on page 711](#)). In the Group Editor, use the Group Page text field to specify URLs to reference pages for the group as a whole. Use the **Group Children's Page** field to specify URLs to reference pages for the individual items within the group. Member URLs can be in the form of templates that use the names of [Data Fields](#) to query for particular files.

Note also that all the content formerly available through the feature called "Vendor Pages" continues to be available from Reference Pages.

Reports

Reports are a resource that provide captured views or analyses of information that can be viewed in the ArcSight Console in PDF, Excel, Comma Separated Value (csv), or Rich Text Format (rtf). You can view the HTML format with your preferred default Web browser.

You can create reports on all events, cases, notifications, and assets that are in storage.

Reports gather data based on [Queries](#) and Trends (["How Trends Work" on page 467](#)), and use report templates to determine display and file formats.

For an overview of reporting, including information on queries, trends, and templates, see ["Building Reports" on page 388](#).

Working with Report Templates, Queries, and Trends

Reports show results of pre-defined queries and trends using custom-designed or provided templates. Once you have source data defined in queries and/or trends, you can design reports to present the data in charts and tables.

The Reports resource includes the following tabs and editors for the following elements that make up reporting:

- Templates (["Using Report Templates" on page 426](#))
- Queries (["Query Settings" on page 321](#))
- Reports (["Defining Report Attributes" on page 394](#))
- Archives (["Archiving and Scheduling Reports" on page 459](#))
- Trends (["Defining Trend Settings" on page 471](#))

A Report Wizard is provided for creating reports quickly. (See ["End-to-End Reporting Examples" on page 430](#) for an example of using the report wizard.) From within the wizard you can choose a data source (one from among available Queries, Trends, [Active Lists](#), or [Session Lists](#)) and one of the available templates to use for the report.

Viewing and Managing Reports

You have the flexibility to broaden or narrow the data extracted from the database using report parameters and conditional logic statements. You can also create delta reports to show the difference between two sets of parameters. With this flexibility, you can create custom reports that are tailored to meet your reporting needs. Also, there are display groupings of "Report Definitions," where you define what report you want to generate, and "Report Output," where the actual reports are generated and stored.

Archived Reports

Once a report is created, it can be saved (archived). Archived reports are retrieved for immediate viewing, without requiring you to regenerate the report. In addition, you can schedule a report for automatic archiving, on a yearly, monthly, weekly, daily, or hourly basis. All reports are displayed at the ArcSight Console in the Report Viewer.

Report Groups

Reports can be created and edited in the <user ID>'s Reports group or the Public Reports group on the Report Definitions tab. Reports can be deleted based on your permissions. Reports created in the <user ID>'s Reports group is available only to that user and those to whom the user gives inspect or edit permission. Reports created in the Public Reports group are available for all users to create, edit, or delete. Reports can then be run from the Reports resource tree and the Viewer panel. For more information, see ["Running and Managing Reports" on page 446](#).

You can manage reports in the Reports window of the ArcSight Console. The Reports resource tree has two tabs: Report Definitions where reports are managed and Report Output where archived reports are stored for viewing. The Report Definitions tab lists and organizes all reports in one of the following groups.

Report Groups	Description
<user_ID>'s Reports	Reports the user has created.
Shared Reports	Reports that other users have already shared with the logged-in user.

Report Groups	Description
ArcSight Reports	Reports provided as defaults, which you can use as-is or to create custom reports.
Downloads	Reports downloaded from various use cases.
Public Reports	Reports to which all users have read permission.

If you have Administrator access you will have another group named All Reports that contains all user report groups and their reports.

The Reports Output tab lists all archived reports. Archived reports are listed as a file on the Report Output tab for quick access and retrieval. When you archive a report from the Report Definitions tab, that report is sent to the Report Output tab. If other users archive a report and share it, the report is listed in the Shared group on the Report Output tab.

Delta Reports

A delta report is one that shows the difference between two sets of parameters used in a single report. The report also shows the data for each of the parameters.

When you run or archive a delta report, an internal event is sent to the Manager. This event contains the following data fields and values.

Delta Report Field	Description
Event Name	Delta Report Generated (Report: <i><ReportName></i>), where <i><ReportName></i> is the name of the report.

Rules can be created using the delta report data fields.

Report Parameters

For date parameters, type in the text fields, click the drop-down arrows or click the time buttons to select a time range.

For date and time data fields, you can also type an actual date value, such as 10/12/2015 8:54:00 AM, or you can use special system variables such as:

- `$CurrentDateTime`: for the date and time the report is run, the system variable is replaced by the current date and time value when the report is run.
- `$CurrentDate`: for the date the report is run, the system variable is replaced with the date value, truncating the time of the day to 0, when the report is scheduled or run.

You can also specify certain date operations with these system variables to add or subtract a number of specified days or hours. For example, you could type: `$CurrentDate - 7d` for seven days before the date the report is run, the condition evaluates to a date which is the current

date minus seven days, or `$CurrentDateTime - 12h`, which evaluates to the current date time minus 12 hours.

Select a **Report File Format** from the drop-down menu.

Reports can be archived in PDF, HTML, Excel, Comma Separated Value (csv), or Rich Text Format (rtf). The default PDF format should be used when archiving reports. Compared to PDF reports, other reports may lose formatting information and will appear differently. In addition, Excel format is more memory intensive than PDF.

ArcSight-Provided Reports

There are more than 200 reports listed in the ArcSight Reports group which you can use to immediately generate reports or use as templates to create or customize your own reports.

The ArcSight Reports group contains all reports further subcategorized in one of these report subgroups.

ArcSight Reports

Report Group	Description
By Attribute	Provides various reports providing details based on connector type, severity, device, event name, source, target, and target port.
By Event Direction	Provides reports based on inbound or outbound attack and alert direction.
By Event Volume	Provides reports based on most frequent and least frequently occurring events.
By Sensor	Provides report by type of sensor or connector, for example, firewall, router, intrusion detection, and so forth.
By System Object	Provides reports by type of object, for example, active lists, assets, cases, and notification.
Custom Reports	Provides varied example reports demonstrating moving averages and vulnerability.
Example Canned Reports	Provides reports showing various types of asset reports and different report layout designs.
Internal Reports	Provides various reports providing information on system usage, performance, rule operation, and resources.

You can run each of the reports to see the type of detail and information each provides. In addition, you can view the notes, settings, and conditions set on each tab to see how each report is constructed. To use a report as a template for creating or customizing your own reports, copy an existing report to a new report group (using drag and drop to copy an existing report to the new report group). You can then rename the new report and start making changes to the new report.

Report Templates

Report templates are a component of Reporting resource tools.

To provide more flexibility in reporting, the report template tools include a rich offering of ready-made templates and a template design wizard for more customized [Reports](#). Template definitions determine how data from [Queries](#) and Trends ("[How Trends Work](#)" on page 467) is displayed in a report. You can create and adjust templates to specify which data is displayed, what visual elements are used (variations on tables, charts, graphs, and so on), the layout of those elements, the report output file format, and much more. A template consists of report design elements, such as headers, footers, title bars, charts, and tables, arranged on a page according to a layout specification.

Templates can accommodate input from multiple queries and show multiple visual elements, such as three charts and a table each pulling from a different data source, in a single report.

For more information on templates, see "[Using Report Templates](#)" on page 426.

See "[Building Reports](#)" on page 388 for an overview of all reporting tasks and tools.

Resources

The Manager handles the logic used to process events as objects called **resources**. Active channels, data monitors, filters, cases, assets, queries, trends reports, rules, and packages are all examples of resources.

A resource defines the properties, values, and relationships used to configure the functions the Manager performs. Resources can also be the output of such a configuration (such as archived reports, or Threat Detector snapshots and patterns).

The Manager has over 30 different types of resources and comes with hundreds of these resources already configured to give you functionality as soon as the product is installed. These resources are presented in the Navigator panel of the ArcSight Console.

This topic provides an overview on working with resources in the Console. Resources in general are discussed in more detail in the topic "Resources" in *ESM 101*.

Resource Attributes

The processed [Resources](#) are composed of several attributes, each of which is a data field with its own characteristics. The data fields common to all resources are described below.

Each attribute has both a **Label** that you see in the ArcSight Console and a unique **Script Alias** you use to refer to the attribute in filters, rules, or Velocity templates. The **Data Type** lets you know how to handle the attribute. (See also ["Resources" on the previous page](#) for information on locked and unlocked resources, and ["Common Resource Attribute Fields" on page 725](#) for information on viewing and/or editing these fields in resource editors.)

Resource Attributes

Group	Label	Script Alias	Data Type	Description
<i>Resource Type</i>	Name	name	String	Top-level categorization of the resource as shown on the Navigator panel, for example, Active Channel, Asset, Rule, and so on.
Common	Resource ID			Read-only field that shows the system resource ID.
	External ID	externalId	String	An identification string suitable for, and which can be referenced by external systems. Common applications of External IDs include appropriate naming for Case and Asset resources that are tracked in common with defect reporting or vulnerability-management systems. Your administrator can advise you on the correct values for this field, if applicable. For Vulnerability resources, this field will be filled in with an ID of the format <standards body> <id>, such as CVE CVE-1999-200.
	Alias	alias	String	An identification string suitable for referencing resources. A given alias will appear in place of the resource's name everywhere it may be seen. Your administrator can advise you on the correct values for this field, if applicable.
	Description	description	String	An editable text description of the resource or other related information. This text appears as a tooltip to any user who has ArcSight Console access to the resource.
	Version ID			A string showing a globally unique version ID for the resource.

Resource Attributes, continued

Group	Label	Script Alias	Data Type	Description
	Deprecated			Indicates whether a resource is current or obsolete. If this field is blank, the resource is current. If this field is check marked, the resource is "deprecated" or obsolete. Click the box to toggle the checkmark on or off.
Assign	Owner	owner	String	One or more users who are interested in this resource.
	Notification Groups	notificationGroups	String	The ArcSight user groups selected from the Users resource tree who should be notified about this resource.
Parent Group	groupNameLink		Resource Group	Each resource group containing this resource. A resource exists in more than one group when you choose Link instead of Copy or Move.
Creation Information	Created By	userName	User	The identity of the user who created this resource.
	Creation Time	creationTime	DateTime	The time that the resource was created.
	Time Since Creation	timeSinceCreation	String	The elapsed time, in days, hours, minutes, and seconds since this resource was created.
Last Update Information	Last Updated By	lastUpdatedBy	User	The identity of the user who last updated this resource.
	Last Update Time	lastUpdateTime	DateTime	The time that the resource was last updated.
	Time Since Last Update	timeSinceLastUpdate	String	The elapsed time, in days, hours, minutes, and seconds since this resource was last updated.

Rules

An ESM rule is a programmed procedure that attempts to correlate incoming network [Events](#) and generates new events that report on correlation when it occurs, as determined by security policy. Rules also apply [Conditions](#) and perform [Rule Actions](#).

Canned rules can be viewed, edited, and used as templates to create your own enterprise-specific or custom rules. To see what's available, browse the description provided with each rule in the ArcSight Console.

Different users can simultaneously create rules from their ArcSight Consoles. Once created, all rules are sent to the Manager, which updates any other individual Consoles. Updates to [Resources](#), including rules, are automatically refreshed every few seconds so that clients get the latest changes from other clients.

Information on creating, deploying, and managing rules is provided in Rule Authoring.

Loading Rules

How you create rules affects the load placed upon the ArcSight Manager. This load is a function of how many partial and full matches are generated by those rules. Since partial matches occur when any condition of a rule is met and full matches occur once all conditions of a rule have been met, poorly written rules can generate many partial matches without generating any full matches.

Also, poorly written rules can generate, in a worst case scenario, one additional event for every incoming event. However, well-written rules have conditions that are restrictive enough to limit partial matches to those events that are likely to participate in a full match. Such rules are also likely to generate very meaningful derived events and they also impose a smaller load on the ArcSight Manager. Therefore it is very important that you carefully plan, write, and test all your custom rules.

See "[Automatically Disabled Rules](#)" [below](#) for more information.

Automatically Disabled Rules

A rule can be manually disabled by an administrator or automatically disabled by ESM. ESM automatically disables improperly written rules that would produce excessive or meaningless events.

The Rules resource tree on the Navigator panel displays a manually-disabled rule as greyed out (). An auto-disabled rule is displayed with a special icon (). It shows with the same disabled symbol overlaid by an ArcSight logo to indicate that the system disabled it.

When a rule is disabled, ESM generates an audit event indicating that this happened so that administrators can follow up as needed. See "[Rule Activations](#)" [on page 871](#) for more information on related audit events.



Tip: About the Rules Status dashboard:

ESM profiles rule performance by measuring their evaluation time on a sampling basis. You can view these results from the All Dashboards\ArcSight Administration\ESM\System Health\Rules\Rules Status dashboard, which includes a collection of data monitors reporting on different rules statistics. Based on information from this dashboard, manually disable rules which you deem expensive.

The Sortable Rule Stats data monitor on this dashboard does not include pre-persistence rules.

Why Rules are Automatically Disabled

Cause	Description
Rule is invalid	An invalid rule is automatically disabled and displayed as broken  in the Navigator. If an administrator configures a rule or related resource in a way that “breaks” the rule and leaves it in an invalid state, the system automatically disables the rule. If a rule is disabled automatically due to an invalid configuration, an Invalid Reason field is displayed in the Rule Editor on the Inspect/Edit panel. When the rule is reconfigured to a valid state and enabled, the Invalid Reason field is no longer displayed. The Invalid Reason field is not displayed for rules that are manually disabled.
Rule is recursive	Rules that trigger themselves in a recursive loop is automatically disabled <i>temporarily</i>. A rule that is automatically disabled due to recursion is re-activated after a time frame that matches the aggregation time frame for the rule. (The default aggregation time frame is 2 minutes.) A rule can be inherently recursive due to a flaw in its design, or temporarily recursive because of some particular events involved. In the first case, temporarily disabling the rule often clears out the problem, and allows the rule to run normally when it is re-activated. If the rule is inherently recursive, it is continuously re-enabled and auto-disabled. The solution in this case is to redefine the rule logic and redeploy it, since it is effectively a “broken” rule.
Excessive event alias matching	This is the number of events matching that alias, independent of other defined aliases. The default limit for event matching is 100000.
Partial event matching	If more than one event alias is defined in the rule, partial matching is the number of events matching the aliases defined before the current one, and for the current one, and for their join condition (if present). The default limit for partial matches of any event aliases is 100000.
Generated event counts	This is the number of correlation events generated. The default limit is five correlation events for each base event the rule processes.
Base event counts	The number of base events used by the rule to generate correlation events.

Why Rules are Automatically Disabled, continued


Cause	Description
Time unit counts	This is the number of time units (minutes) that passed since the current rule activated. The default is 1000 correlation events in one time unit.
Number of rule triggers exceeds configured limits	Number of rule triggers exceeds configured limit of 1000 firings per minute for the same aggregated values. A rule that exceeds configured limits show as disabled (🚫) in the Navigator, and offer a right-click option for the user to manually disable it permanently. To change this setting, do so in the <code>server.properties</code> file. For example, if you want the limit to be 10000 instead of 1000, enter this setting: <pre>rules.max.fan-out.time-unit.ratio=10000</pre> For information on how to set properties, see "Managing and Changing Properties File Settings" in the ESM Administrator's Guide. Note: A rule in this state continues to attempt to run until the user disables it permanently by right-clicking it in the Navigator and choosing Disable.
CPU usage has exceeded threshold	ESM takes the aggregated evaluation time of all deployed rules. If a rule's evaluation time exceeds 50% of this aggregated time, the rule is automatically disabled. To change this setting, do so in the <code>server.properties</code> file. For example, if you want 60% instead of 50%, enter this setting: <pre>rules.max.fractional.cpu=60</pre> For information on how to set properties, see "Managing and Changing Properties File Settings" in the ESM Administrator's Guide.

For rules that are disabled automatically, right-click the disabled rule and select **Disable** so that the rule is permanently disabled until you can fix the rule. If you don't manually disable these rules, they continuously attempt to run, then are enabled and disabled by the system in a cyclical manner. This can impact system performance.


Rules Processing and Correlation

A rule has three parts: a condition, threshold and time window aggregation, and an action. The condition states if exists and satisfies expressions and the action states do expressions. A rule states if [one or more conditions] exist and satisfy the rule, then do [action expressions]. A rule can have one or more rule conditions. If there is one condition, the rule acts as a filtering tool. If there is more than one condition, the rule acts as a correlation tool. A rule can be created for any incoming event from one or more event generators, with various conditions, logic statements, and threshold and time window qualification of events.

Components of a Rule




The Correlation Engine, a sub-component of the Manager that handles rules, is not the same as a database query engine. For example, the Correlation Engine can perform a complex join across several events in real-time and aggregate the response to these events. In order for the Correlation Engine to do this in an efficient manner, it keeps a list of events that match each condition. These are referred to as **partial matches** because they satisfy part, but not all, of the rule's conditions. As new partial matches occur, the Correlation Engine attempts to pair them with previous partial matches in order to construct a full match. At that point the Correlation Engine may aggregate that match with others while it waits to pass some threshold (which can be either time or a target number of full matches). If the threshold is passed the Correlation Engine generates a derived event and performs the other actions associated with the rule.



It is important to note that all rules containing a specified threshold and a time window expiration follow a certain process in order to generate a derived event and perform an action. If a rule's threshold is passed, but the time window expiration has not been met, then the Correlation Engine compensates for this by generating a derived event, performing an action, and moving (or sliding) the time window until it expires. If this rule process was not in place, under certain conditions, rules would trigger on nearly every event in a short amount of time and which would cause a large amount of useless events to be displayed or actions taken.

For example, assume that you created a rule with an event threshold of 4 and a time expiration window of 3 minutes that sends a notification every time the threshold is met. This rule's process would look like the following:



In this example, the 4th incoming event occurred before the time window expired, so the rule triggered at the 4th event and the time window shifted adding another 3 minutes. Within the 2nd 3-minute interval, the rule restarted its incoming event count; however a 4th event did not occur so the rule did not fire. Note that the time window did not expire until the 5th minute had passed. If a 4th event had occurred before that time then the shifting process would have begun again. If you were to show the rule chain for this example, it would display the information for incoming events 1-4 that occurred within the first 3 minutes. Time windows expiration triggers fire at the minute boundary unless the next time window starts before the minute boundary.

The Rules resource tree in the Navigator panel offers a default collection of rules that you can use directly or as a template for creating your own custom rules.

For example, there are rules predefined to detect and perform actions based on system rules processing and SmartConnector status. Other rule groups detect and respond to attacks and suspicious activity, specific types of attacks on various sensor types, network components, or assets, and report attack results or successes.

Rule Groups

Rules are organized into groups to store similar rules in one location. The Rules tree in the Navigator panel organizes rules into the following groups.

Rule Groups	Description
<userID>'s Rules	The user's home directory, where they have read/write permissions to author rules.
Shared Rules	Rules that establish the permissions for the current user.
Real-time Rules	Rules that are run against real-time events.

Rule Groups	Description
Public Rules	The rules that any user can read.
System Rules	The global rules provided by ArcSight.
Unassigned	Rules that do not belong to any directory. These can be rules that have not been inserted into any directory, or their parent directory has been deleted.

If you have Administrator access you will have another group named All Rules that contains all user rule groups and their rules.

Scheduled Rules

You can deploy scheduled rules to run at a specified time interval (such as hourly, daily, or monthly). This is a useful alternative to real-time rules in situations where you want to deploy rules that take into account historical data along with live data, or when you simply want to control when the rules are run. The scheduled rules engine can process historical data, take real actions, and generate correlated events which are the same as those generated by the real-time rules engine.

Only rule groups can be scheduled. To schedule one or more rules, you add the rules to a rule group, and then edit the rule group to define a job schedule. For more information, see ["Scheduling Rules" on page 558](#).

Rule-triggering Timing

Rule-processing sessions are associated with Group By tuples (for example, a particular pairing of source and target address).

A match occurs when all the conditions of the rule are met.

The first match associated with a new tuple creates a new session. It also triggers onFirstEvent and an OnEveryEvent. The system then sets the start time for the first time window.

Subsequent matches will trigger onSubsequentEvents and onEveryEvent.

If enough matches occur to pass the threshold count **before** the time window expires (which is defined as $\text{start time} + \text{time window} > \text{current time}$), then the Manager triggers onEveryThreshold and one of either onFirstThreshold or onSubsequentThreshold, then resets the start time for the next time window.

If a time window ends without meeting the threshold, then final aggregation occurs. The onTimeWindowExpiration option is triggered and the session is disassociated from the tuple.

The next match with the same or a new tuple will cause the whole process to repeat.

Rule Chains

Sometimes you want to capture or act upon a series of correlated events within a specified interval or at a particular threshold. When rules are designed to trigger in a series, they are referred to as rule chains.

Variables

You can use all of the dynamic time parameters you see in the Active Channel Editor and elsewhere, such as `$Now` and `$CurrentDateTime`. The same is true for time elements, including `s` (second), `m` (minute), `d` (date), `M` (month), `w` (week), and `y` (year). To use any event data field as a variable, express its displayed name as a one-word "camel cap" string prefixed with a dollar sign; for example, "Source Address" would be `$sourceAddress`.

Rule Actions

Rule actions are automatic procedures that occur when all rule [Conditions](#) and threshold settings have been met. (See also ["Rules" on page 1034](#).) You can choose to be notified of a triggered rule at the ArcSight Console or through the Notifier (see ["Notifications" on page 1010](#)), have information about the [Events](#) that triggered the rule sent to a case or active list (see ["Cases" on page 877](#) or ["Active Lists" on page 820](#)), or automatically execute a command line function. You can also assign more than one rule action to any rule.

The task steps for these activities are available in ["Managing Rule Actions" on page 532](#).

Rule Conditions

A rule is a programmed procedure that can analyze network [Events](#) and generate additional correlation events, as determined by security policy. (See also ["Rules" on page 1034](#).) When creating rules, you define the rule events and [Conditions](#), thresholds, and [Rule Actions](#). Conditions define which events trigger the rule, thresholds set when a correlation event is generated, and actions state which responses are taken when a correlation event is generated. To define rule events and conditions, thresholds, and actions, begin by determining the following:

- Which event occurrences do I want to be aware of? This determines the rule's **events** and **conditions**.

- How many times do I want the event or events to occur and within what time frame? This determines the rule's **threshold**.
- What actions should automatically occur when an event is generated? When should those actions occur? This determines the rule's **actions**.

A rule requires at least one event and one condition. When you create or edit a rule, the ArcSight Console provides a Conditions tab in which you can specify events and define the conditions for a rule. (The Conditions tab is described in the topic, "[Common Conditions Editor \(CCE\)](#)" on page 887.)

Rules are first constructed by creating condition statements. Condition statements contain a data field, logic operator, and data field value; so you can create complex logical expressions by combining one or more individual conditions to match the events you want to trigger a rule.

When you first create a new rule, a default event named event1 appears as a branch under the Event conditions tree for the new rule. (The event name is also commonly referred to as the event alias.) You can use this name or select a different event to use in the condition. Since rules can have numerous events, event names should be unique and descriptive within the same rule. For example, if monitoring Cisco Router denied events, Cisco Router denied could be the event name. The event name appears as a branch under the Event conditions tree.

When defining the condition for an event, the Conditions tab provides three columns, Name, Operator, and Condition. These three columns are combined to create <data field> <logic operator> <data field value> condition statements. For example, if monitoring a Cisco Router, the condition statement could be Device Product = Cisco Router: Device Product as the data field, = as the logic operator, and Cisco Router as the data field value.

When adding conditions, you need to decide how to tie the new condition to any existing conditions. To add more condition statements to an event, you can use logical operators AND, OR, or NOT to specify how to evaluate the condition statement that contains more than one individual condition.

Besides specifying events in a condition, you can also add filters, assets, and vulnerabilities to rules as new conditions. A filter condition monitors if an event occurs in a particular filter. If an event does occur in that filter, a correlation event is generated (see "[Specifying Rule Conditions](#)" on page 513). Asset conditions state whether your enterprise assets are targets or sources of events. An asset condition states if an event occurs and the selected asset is the source or target, generate a correlation event. Finally, you can also use an existing enterprise vulnerability to create a rule condition. A vulnerability condition states if an event occurs with the vulnerability selected, generate a correlation event. For more information on vulnerabilities, see "[Modeling the Network](#)" on page 119.

In some cases, however, you may want to specify more complex rule processing to restrict the events that actually cause a rule to fire. There are two additional elements you can include to specify more complex rule conditions: rule thresholds and aggregation. See ["Specifying Rule Thresholds and Aggregation" on page 528](#).

Rules Editor

The Rules Editor is a panel in the ArcSight Console for creating and editing rules.


The rules you create or edit are stored in .ARL (ArcSight Rules Language) files.

For more information, see ["Rules Authoring" on page 507](#).

Saved Searches

You create and save searches only in the ArcSight Command Center. After you save these searches, they appear on the Console as resources.

With these resources, you can leverage the grouping and packaging functionality available in the Console.



For related information, see:

- ["Managing Resources" on page 711](#)
- ["Managing Packages" on page 730](#)
- The topic, "Saved Queries (Search Filters and Saved Searches)" in the Command Center User's Guide

Schema

The schema is more than 400 data fields of the normalized data recorded by the device (sensor) that reports events to the SmartConnector. The schema is the culmination of the normalization process, and the backbone of the data structure that drives correlation.

The schema also includes fields that support resources that operate on other resources, for example, actors, assets, and cases.

The schema can now also be expanded with user-defined fields. Global variables enable you to define a variable that derives data from fields in the schema, which can be used in multiple places (see ["Global Variables" on page 609](#)).

Avoiding Field Naming Collisions

With the addition of user-defined fields to the schema comes the possibility of name collisions. In most cases, field names, regardless of type, must be unique to be resolved. The following attributes are checked to verify that names are unique for all types of data fields:

Fields Validated for Uniqueness

Field Type	Field Validated
Event	name, alias, field name, display field name
Actor	name, alias, field name, display field name
Asset	name, alias, field name, display field name
Case	name, alias, field name, display field name
Custom columns (public)	N/A
Custom columns (private)	N/A
Global variable	name, alias, field name
Local variable	name, alias, field name
Domain field	name, alias, field name

The Manager uses the following policy to manage potential naming collisions.

- The Manager grants names on a first-come-first-served basis. The domain field or global variable that comes later with the same name as another field will either be marked as 'disabled' if added in batch mode (such as from an archive file or package) or 'denied' when being created directly from the ArcSight Console.
- Name collision is allowed among resource and event-based system fields. For example, the name field of event can be the same as the name field of an actor.
- Global variable names must be unique across all types of schema fields. For example, a global variable cannot have the same name as a domain field or event.
- The name of a local variable must be unique across all types of fields: event fields, resource-based fields, global variables, and other local variables in the same containing resource.

When Requesting a Name for a ...	Can the name be the same when in use by		
	Global Variable	Local Variable	Event, Actor, Asset, Case
Custom Cell	Yes	Yes	Yes
Domain Field	No	Yes	No
Global Variable	No	Yes	No
Local Variable	No	No	No

The following exceptions apply to avoid naming collisions with existing customer-created fields and ArcSight-supplied global variables during upgrade to a future release:

- Existing custom columns added to active channels (see ["Customizing Columns" on page 250](#)) are excluded from name collision validation. Custom columns can have the same names as event fields, resource fields, global variables, and local variables, and vice versa.
- New global variables can have the same name as an existing local variable. A new local variable cannot have the same name as an existing global variable, but if a local variable already exists with a particular name, a global variable with that same name can be added without a name collision error.

Event Fields

- Most information reported by sensor devices are in the main event fields.

The information is accessible from Rules, Filters, and Reports. The events from a supported sensor might include three different fields - encryption failure, encryption success, and error - that all contain messages. These three are all mapped to the 'message' field.

- Usage clarification for many fields.

This information to help you write rules or data monitors that use the variety of possible values.

Precise Event Categorization

ESM categorizes events across six dimensions:

- the object acted on by the event
- the action represented by the event
- the technique used to achieve the action
- whether or not the action succeeded
- the security significance of the event
- the class of device that reported the event.

See ["Categories" on page 878](#). This scheme supports focused rule authoring and data monitor construction.

Events describe either an action or a state. Actions are attempted against a particular object and may succeed or fail. There may be many ways to attempt a particular action against an object (such as different ways to exploit an exposed vulnerability). States describe the status of a particular object, and these states may be known to be true -- or they may be hearsay. Events all have some significance to the security profile of the protected network. Finally, it is interesting to know what sort of device is reporting the event.

If we look at Snort SID 103, we discover that it is a report of a scan searching for pre-installed subseven 22 back doors. We would categorize these events as follows:

Security Significance	/Recon
Behavior	/Communicate/Query
Technique	/Scan/Service
Device Type	/IDS/Network
Outcome	/Attempt
Object	/Host/Application/Backdoor

In this case, a network intrusion detection system (IDS) would observe an attempt to communicate with a backdoor and infer that this was part of a service scan attempting to discover pre-installed instances of that backdoor. Naturally, this implies an external connector is performing reconnaissance on the protected network.

Search Filters

You create and save search filters only in the ArcSight Command Center. After you save these search filters, they appear on the Console as resources.

With these resources, you can leverage the grouping and packaging functionality available in the Console.

  **Caution:** Do not exceed more than 10,000 resources in a group.

For related information, see:

- ["Managing Resources" on page 711](#)
- ["Managing Packages" on page 730](#)

- The topic, "Saved Queries (Search Filters and Saved Searches" in the *ArcSight Command Center User's Guide*
- The topic, "Content Management," in the *ArcSight Command Center User's Guide*

Send Logs

Several system components provide various types of information to log files. For example, the Manager logs are located in: `ARCSIGHT_HOME/logs/default/server.log`. Various Manager utilities write logging information into different sets of log files. (The archive utility writes to the `archive.log`, the database init utility writes to the `dbwizard.log`, and so forth). Each of those sets can consist of multiple files. The number and size of the log files are configurable on the Manager under `ARCSIGHT_HOME/config/server.properties`.

ArcSight Console and connectors also generate and store log files.

Software Support may request log files and other diagnostic information to troubleshoot problems. The **sendlogs** utility automatically locates log files and compresses them.

The Send Logs utility does not send any log files automatically; you will be instructed by the support representative on how to send the log files.

Guidelines for Using the Send Logs Utility

- Although Send Logs is accessible to any user logged into a component, only administrators have permission to use the Send Logs utility.
- SmartConnectors must be running version 4037 or newer to support remote collection of logs with Send Logs (using an ArcSight Console or the Manager)
- You can only collect local logs on SmartConnectors or ArcSight Database. That is, if you run the Send Logs utility on ArcSight Database, only the database log files are gathered.
- You can run the `sendlogs` utility on a component even when the component service is down. If the Manager is down, you can only collect its local logs. However, if you need to collect the database logs as well, use the `arcdt` command on the Manager.
- All log files for a component are gathered and compressed. That is, you cannot select a subset of log files that the utility should process.
- The `sendlogs` utility generates a compressed file on your local system. The compressed log file is created in `ARCSIGHT_HOME/tmp/logs/<Filename>.zip`.
- You can review the compressed file to ensure that it contains only the desired and appropriate amount of information.

- You can remove or sanitize information such as IP addresses, host names, and e-mail addresses from the log files before compressing them. The options are to (a) send logs as generated, (b) remove only IP address, or (c) remove IP address, host names, and e-mail addresses.

Options for Running Diagnostics and Sending Logs

There are two ways to launch the **Send Logs** utility:

- As a wizard from the ArcSight Console Tools menu. (See ["Using the Network Tools" on page 73.](#))

For details on using the sendlogs utility from the command line, see the "Administrative Commands" section of the Administrator's Guide.

- From the command-line interface of each component, using the command **arcsight sendlogs** from ARCSIGHT_HOME/bin on the ArcSight Console or Manager.

For details on using the sendlogs utility from the ArcSight Console, see the "Configuration" Chapter of the *ESM Administrator's Guide*. It is under "Configuring Manager Logging" then "Gathering logs and diagnostic information."

You can also use the **arcctl** command to run specific diagnostic utilities from the Manager command line.

Starting the Send Logs Wizard on the ArcSight Console

You can launch the Send Logs Wizard from the ArcSight Console by choosing **Tools > SendLogs** from the Console menu.

The first time you run the utility from the Console, you are prompted to select the components from which to gather logs and diagnostics. Some additional first-time settings are also required (such as notification details, time ranges, and options to gather diagnostics for session wait times, thread dumps, and database alert logs). The wizard remembers most of these, so that on subsequent runs you can choose to use to retain the original configuration.

From now on when you start the Send Logs wizard, you will get this dialog. If you want to retain your original settings, select **Use current settings to gather logs**. If you want to re-set the configuration, select the **Change/Review settings before gathering logs**. option. From either of these initial dialogs, the wizard guides you through the process of collecting logs and diagnostic information.

For a full description of all options and settings on the Send Logs utility, see "Configuring Manager Logging" in the *ESM Administrator's Guide*.

Session Correlation

A session is information about the actors behind your network traffic that applies for a limited and specific period of time. Session information can be used to answer questions such as: "Who is in the New York office?" or "How many people are in meetings?" or "Are users accessing this resource according to company policy?"

Session correlation is a set of tools that capture session information to not only identify the assets involved in network traffic, but also the users, or actors, behind the traffic. (See ["Understanding Session Correlation" on page 565.](#))

Session correlation makes it possible to map users to assets at specific time periods. This is especially valuable for identifying who is doing what on your network from which assets and when, especially when the asset IDs themselves may be variable (such as DHCP or VPN logins).

Why Session Correlation Matters

Monitoring traffic on your network generally means processing data about the assets involved in the network traffic. However, there are times when asset data alone is not sufficient to detect potential threats to your network.

For example, users who log into the network on VPN or DHCP connections are assigned different IP addresses every time they log in. When sensors report events to SmartConnectors, they are only identified by their assigned IP address, which means that you may be missing a whole spectrum of activity from mobile assets, such as laptops and PDAs and remote offices.

Whether accessing your network by using assets with fixed or variable IDs, it is often the user (the "actor") involved in the network activity whose actions you want to correlate with other event data. This enables you to track who is doing what on your network and when, and what they are doing in subsequent log-in sessions.

Capturing data about who is involved in network traffic as well as what assets are involved also adds crucial verification data to your correlation process. For example, three failed login attempts from a particular IP address can trigger a rule. But if that IP address is assigned to three different assets in the timeframe evaluated, session correlation makes it possible to clarify that the three failed login attempts were not executed by the same user.

Session Lists

Session Lists are similar to [Active Lists](#), with the following major differences:

- Session Lists always have Start Time, End Time, and Creation Time fields.
- Session Lists partition data into weekly partitions because the lists can grow very large over a period of time.
- Session Lists do not have to fit entirely in memory.
- Session Lists are optimized for efficient time-based queries.

Session Lists can monitor activity based on any [Rules](#)-driven combination of [Events](#) attributes or set of custom fields. For example, session lists are very useful for tracking suspicious or hostile IP addresses as well as targets of attacks that may be compromised.

While you can populate session lists "manually" (adding entries from grid views or the Session List Editor), you should use session lists in conjunction with rules specifically tailored to work with them. Rules can dynamically add and remove entries on lists, thereby making them a flexible information-gathering tool.

You can open and edit session lists in grid views.

Session lists are not continuously re-evaluated and are not time-window constrained. Session lists draw from the event stream on the basis of their event or field/rule definitions and any rules designed to affect them.



Caution: Be careful about using large session lists in filters. This may severely impact system performance.

In addition to their integral definitions, you can apply temporary (not saved) filters to session list grid views. Click the status description in the **Filter** line in the view header to use the [Common Conditions Editor \(CCE\)](#).

Use the set of default items in the Session Lists resource tree for templates or for operational monitoring with minor modifications. For example, use the ArcSight User Sessions list to watch activity related to logins.

If you have Administrator access you will have another group named All Session Lists that contains all session list groups and lists.

See also "[Session Correlation](#)" on the previous page.

SmartConnectors

SmartConnectors are collectors of security event information generated by multi-vendor security devices throughout your enterprise. SmartConnectors normalize and correlate this data into [Events](#), expressed as ArcSight Messages, which are forwarded to the Connector Data Manager (a component of the [Manager](#)) for further processing. SmartConnectors can reside on

a device, on the Manager, or on a host machine. (For more information on SmartConnectors, see also ["Managing SmartConnectors" on page 175.](#))



Caution: Do not delete a Connector resource at the Console unless the corresponding SmartConnector is first uninstalled from the device it is running on. If the SmartConnector running on the device has not been uninstalled, and its Connector resource is deleted, the SmartConnector will lose its connection to the Manager, causing the SmartConnector to start caching events and eventually dropping them.

Related topics:

- ["Operational Status" below](#)
- ["Configuration" below](#)
- ["Zones" on the next page](#)
- ["Upgrading" on page 1053](#)
- ["Filtering" on page 1053](#)

Operational Status

SmartConnectors display their operating status conditions next to their names in the Connectors resource tree in the Navigator panel.

SmartConnector Status Conditions

Status Condition	Description
running	The SmartConnector is operating normally.
down	The SmartConnector is not connected to the ArcSight Console, therefore no events are being received.
stopped	The SmartConnector is responding to commands sent from the Console, but events aren't being received.
paused	The SmartConnector is responding to commands sent from the Console, but events aren't being transferred and are remaining in the SmartConnector's cache.

Configuration

You can configure SmartConnectors to set a specific priority level for events that match specific criteria. One of the typical applications of this is to change the default priority mapping. By default, SmartConnectors will map the device priority (which may contain multiple levels) to the standard priority levels: **Very-High**, **High**, **Medium**, and **Low**. For example, if a device has eight priority levels (0-7) where 0 is the highest priority, then most likely 0 and 1 will be

mapped to Very-High, 2 and 3 to High, 4 and 5 to Medium, and 6 and 7 to Low. You can use this feature to change this behavior and make the SmartConnector set the priority based on different parameters. For example, assume two firewalls, one of which is your production firewall and the other an internal firewall used for testing. You can configure the SmartConnectors to set the priority to Low for all the events coming from the internal firewall and leave the rest of the events with the default priority mapping.

SmartConnectors can be configured to optimize their performance and increase their functionality. SmartConnectors can be configured to enable aggregation, batching, and time zone correction functionality. You can also send control commands from the ArcSight Console to SmartConnectors to manage the flow of events.



Note: SmartConnector configuration also affects the ability to automatically create the assets that represent network devices. Each SmartConnector needs to report an IP address or hostname for its sensor so its events can be identified on the network. See the configuration guides for your SmartConnectors to ensure they are reporting this information.

For information on how to import and export SmartConnectors configurations, see ["Importing and Exporting SmartConnector Configurations" on page 205](#).



Tip: SmartConnectors can send event information to the Manager in a compressed format using HTTP compression. Using compression lowers the overall network bandwidth used by SmartConnectors dramatically, without impacting their performance. By default, all SmartConnectors have compression enabled. You can disable compression on SmartConnectors by modifying the `ARCSIGHT_HOME/user/agent/agent.properties` file as described in "Disabling Event Compression" in "Configuring ArcSight SmartConnectors" in *SmartConnector User's Guide*.

Zones

Network zones are address-based network zone information as reported by or assigned to connectors and integrated as an asset property. You can access zone through the [Zones Tab](#) of the Assets resource trees, and the Zone Editor.

The system can gather and integrate zone information by any of the methods. Only one method can apply with a given connector.

- If an AUP file (ArcSight SmartConnector [Content](#) update) is installed with a connector, the zone information, if present, provides addressed-based recognition.
- If a `defaultzones.csv` file is installed in the connector's `ARCSIGHT_HOME/system/agent/acp` directory, it overrides an AUP file if present.
- If the various zone URI values are set in the Connector Editor, in the Network section of the Networks: Content tab, they override URIs from an AUP file, a `defaultzones.csv` file, or the defaults.

Upgrading

You can centrally manage, configure, and update SmartConnectors remotely. You can use the Upgrade SmartConnector utility on the ArcSight Console to install newer versions of SmartConnector software for managed devices, and to review which versions are currently installed. (See ["Upgrading SmartConnectors" on page 209.](#))

The connector upgrade utility is one of control commands available on SmartConnectors. (See ["Sending Control Commands to SmartConnectors" on page 195.](#))

Filtering

SmartConnectors can also act as a filtering tool between devices and the Manager, using filtering conditions. Filtering conditions are set with a combination of AND or OR statements and data field values. Extraneous events are filtered out to minimize the number of events sent to the Manager and analyzed in the ArcSight Console.



Note: Events filtered out by SmartConnectors are not reported to the Manager, so they won't be stored in or available later from the ArcSight Database.

For information on how to import and export filters on SmartConnectors, see ["Importing and Exporting SmartConnector Configurations" on page 205](#) (especially the topic on ["SmartConnector Filters" on page 206](#)).

SMTP

SMTP (simple mail transfer protocol) is used to send e-mail. An SMTP server must be configured either at install time or through context (right-click) menu e-mail settings. For [Notifications](#), the relevant fields are "from address", which designates the e-mail address of notification e-mail sent, and the "outgoing e-mail server," which is the SMTP server used to send e-mail. It is important to ensure that the "from address" specified is one that will not be rejected by the SMTP server, since some SMTP servers will reject unknown e-mail addresses. POP3 and IMAP can be used to check for e-mail acknowledgments.

You can specify these options at install time, or through context (right-click) menu e-mail option settings. For acknowledgements, the relevant fields are "incoming mail server," which is the POP/IMAP server to specify to check e-mail, "incoming mail protocol," which is either POP3 or IMAP, "account" and "password," which are the login name and password to access the mailbox from the incoming mail server. Note that replying to mails from the notification "from address" should reach the mailbox accessible to the "account" login.

Sortable Field Sets

All fields are indexed, so all fields are sortable. Sortable field sets used to include fields that were indexed. Now they are simply a way to manage access, since sortable field sets are associated with user groups to control access, through the ACL Editor, which edits Access Control Lists. See ["Editing Access Control Lists \(ACLs\)" on page 107](#).

This sorting is represented by the ascending and descending **Sort Column** and **Remove Sort** commands you can apply in the headers of grid view columns. This is also the sorting that you apply through the **Sort Fields** tab in the Active Channel Editor when creating or editing channels.

Enabling all fields for sorting, or allowing on-the-fly sort indexing for previously unindexed fields, are both impractical for real-world performance. The practical solution is to select and index the most order-significant or frequently used fields and to make these fields readily available in clearly marked sets. Therefore, field sets are available from a Navigator panel resource in Active Channels called Sortable Field Sets. (In the Navigator, choose **Active Channels**, and click the **Field Sets** tab.)

Sortable field sets are like other [Field Sets](#), except that they are composed only of fields for which sort indexing has been enabled.

The selection of sortable fields and the named sets these fields are collected in are often customized during initial installation for an enterprise, and are usually tailored further after production use begins. Therefore, a reliable list can't be published in advance.

If you try to add an unsortable field to a sortable field set, or try to select sorting for an unsortable field in the **Sort Fields** tab in the Active Channel Editor, the ArcSight Console alerts you about the field's status.



Caution: Here are reminders for using sorting field sets.

- Sortable fields belong to exclusive sets. This means that if you use a sortable field from one sortable field set to control an active channel, you cannot use sortable fields from other sets as secondary sort controls.
- Users should not edit the field sets in the System Field Sets folder. If edits do occur by mistake, the system will auto-restore those resources to their defaults in about an hour.

[Variables](#) are not subject to indexing and therefore are not candidates for sortable field sets.

Sorting Columns in Grid Views

In grid views (including [Active Channels](#)), the names of fields in column headers are indicated with a double arrow icon  and the **Sort Column** right-click command is enabled. This applies to all fields.

To sort a list per a particular column, right-click over the column name and choose **Sort Column**.

If a field is already sorted, one of two additional icons is shown next to the column name indicating which direction the sort is applied.

- A down arrow  indicates a "top-down" sort is in effect on that field. (For example, when the event End Time field is sorted top-down, newer events are displayed at the top of the list and older events at the bottom. When the Priority field is sorted top-down, events are listed from higher to lower priority.)
- An up arrow  indicates the reverse ("low-to-high") sort in effect on that field. (For example, if the event End Time were sorted this way, older events show at the top of the list and newer events at the end. Similarly, a reverse sort on the Priority field would put low priority events at the top of the list.)

When **Sort Column** is chosen on a sortable field, a "low-to-high" sort is applied first (for example, show events from lowest to highest priority if the Priority field is sorted). If **Sort Column** is selected again, the sort order toggles to the reverse of the previous sort (high to low priority, per our example). The **Remove Sort** option disables the sort and returns the list to its unsorted state with regard to that particular column.

Multiple columns can be sorted simultaneously. The most recently applied sort will take precedence.

See also:

- ["Applying a Field Set to an Active Channel" on page 236](#)
- ["Sorting Events in the Active Channel" on page 238](#)

Threat

The means by which the potential of a threat connector to adversely affect an automated system, facility, or operation, can be manifest. A potential violation of security.

Threat Evaluation

The Manager incorporates a system of security-threat evaluation that culminates in the Priority field you often see in views, reports, or event details. The Priority field uses a scale of 0-10 to rate incoming [Events](#), with 10 being the most-significant value. Naturally, you use Priority field [Threat](#)-evaluation values as a factor in many types of analyses and [Rules](#)-driven reaction or [Notifications](#) scenarios.

Evaluation Process

Threat evaluation is *always on* and applies to all the events received by the Manager. The evaluation process consists of:

1. Identify the targeted asset.

The identification process uses (in this order) the Target Address, Target Host Name, Target MAC Address, or relevant asset address range to classify the targeted asset.

2. Identify the targeted vulnerabilities.

Using the targeted asset as a key, the Manager looks up applicable vulnerabilities.

3. Match the targeted vulnerabilities with the vulnerabilities of the targeted asset.

When matches occur, one is chosen and placed in the Event Vulnerability field.

4. Compute the event's threat-priority value.

It is at this point that ESM performs the computation involving model confidence, relevance, criticality and severity (in this specific order), as described further in the next section.

Evaluation Definitions

The Priority field is a calculated value. It uses a formula that processes the contents of certain [Prioritization Fields](#) that help assess the potential security impact of an event. These fields use information about specific [Assets](#) and [Vulnerabilities](#) to establish models, and a confidence factor concerning the appropriateness of those models. Given confidence about a particular asset/vulnerability model, events directed at that asset can then be evaluated against a combination of factors that include relevance, criticality, and severity.

An event has **relevance** as a threat if it contains an [Attack](#) signature that is genuinely applicable to the targeted device, and the device is in a posture that would permit a successful attack. For example, is the event aimed at a valid port, and when the port was checked, was it open?

An asset's degree of **criticality** is based on the way it serves your enterprise, as seen from the perspective of the network's asset categories. For example, a server could be categorized among your "Very High Criticality Assets" because it handles customer financial transactions.

An event has **severity** if the targeted device is of a more sensitive type that is known to be subject to compromise, and the source of the event has been identified as a hostile or suspicious entity. Specifically, this is the value found in the Device Severity field. For example, did the event originate from an arch competitor on your Hostile List and was it aimed at a router on your Compromised List?

These three factors, when enabled by a suitable model confidence value, are averaged to produce the value that appears in the Priority field. If a suitable model confidence value isn't present, then severity and criticality are averaged to produce a value for Priority. The exploited vulnerability is also recorded in the event's vulnerability field. See ["Viewing an Exploited Vulnerability" on page 233](#).

Maintaining Model Confidence

The asset/vulnerability model confidence for various network devices is based on correlations between the asset and vulnerability resources you can see in the resource trees of the ArcSight Console's Navigator panel. Fresh vulnerability information that correlates well with a particular asset's identification results in greater model confidence.

Stated more directly, the model is the sum of the resources that describe the protected and external networks: assets, asset ranges, asset categories, network zones, and certain active lists.

While asset and vulnerability information can be updated manually, it is more practical to refresh this information by automated means such as vulnerability scanners. (See ["Reporting on Output from Vulnerability Scanners" on page 168](#).) ESM can automatically import vulnerability information from certain scanner products. Information drawn successively from the same scanner product is overwritten when duplicative; information from different products is additive. Information about new assets or vulnerabilities generates new resource references, and the Manager automatically matches the new references to their opposites, whether new or old.

Using Threat Evaluation Information

While the Priority field has many obvious uses, starting with simply sorting the events in grid views, there are other ways to put this and its underlying information to work.

Rules, reports, filters, and any place you can apply logic can use the threat-evaluation operators described in ["Priority Calculations and Ratings" on page 1019](#). You can also use the

values described in ["Prioritization Fields" on page 1018](#) to perform many threat-related functions.

Limitations and Workarounds

Because it is dependent upon a certain amount and type of event data, threat evaluation can be inhibited by the following factors:

- A correlation event, produced by a rule or a data monitor, may not be populated with enough information. Only fields used to group by will be populated in correlated events. Without enough information (such as targeted asset or severity) the threat evaluation will not be able to make a sound decision on the event's priority.
- Over-population of correlated events can also inhibit results. Some rules are only used to maintain active lists. These rules do not generate useful new information, but the group by they need to use in order to collect the information for an active list may give them the appearance of a seriously offensive event.
- Rules offer the option to set your own priority. If a rule populates the priority attribute, then a threat model component will not change that value.

To compensate, you can use these techniques:

- Use the Priority field's value to control when you do and don't notify.
- If a rule is inferring some new piece of information (such as the classic Brute Force Login Attempt), then make sure that you group by sufficient information to be able to characterize the threat later. In the BFLA case, that would mean using the source and target addresses from the base events and setting the severity attribute to, for example, Low; the BFL Success rule, on the other hand, would set severity to Very High.
- If the rule is a bookkeeping rule, try to copy as little information forward as you can, set the severity to low, and set the category to /informational.

Thresholds

There are two types of thresholds: rule thresholds and event thresholds.

A rule threshold is the point at which a rule is triggered and a correlation event generated.

An event threshold is the number of times the event must occur before triggering the rule threshold.

A rule can have a threshold that states when the rule is triggered and also specify a threshold for each rule event. For example, thresholds can be created so that a rule is triggered only after all the events in the rule have occurred a set number of times.

See also ["Rules" on page 1034](#). See also ["Events" on page 995](#) and ["Event Categorization" on page 997](#) for information related to events.

Time Error Correction

In the context of the ArcSight Console, time error correction means the synchronization of time between a network device, its SmartConnector, and the Manager.

See also ["SmartConnectors" on page 1050](#).

Timestamps

Because timestamps are a key element in network security analysis, it is important to clarify the location, source, and context of the timestamps.

All timestamps are stored as Coordinated Universal Time (**UTC**) times.

The ArcSight Console presents timestamps in the local time zone of the host computer using the Java Locale facility.

Log timestamps are produced by the local JVM for that component and are written using the Java Locale facility.

Timestamps are kept in epoch time, an integer value representing the number of seconds since January 1, 1970, at 00:00:01 (UTC). Timestamps cannot be earlier than that date/time. The largest integer (number of seconds) that can be stored for this value limits the timestamp range to January 19, 2038 at 03:14:07 (UTC). No timestamps can be after that date/time.

See also ["Timestamp Variables" on the next page](#).

Timestamps for Security Events

Multiple timestamps are applied to events in the course of processing.

Timestamps for Events	Context
Device Receipt Time	The timestamp applied by the source sensor device upon receipt of the event.
Connector Receipt Time	The timestamp applied by the SmartConnector's JVM (Java Virtual Machine) when the event is received from the originating sensor device.
Manager Receipt Time	The timestamp applied by the Manager's JVM (Java Virtual Machine) when the event is received from the SmartConnector.
Start Time	The time at which the event actually began, as recorded by the source sensor device or, possibly, a secondary source monitored by that device.
End Time	The time at which the event actually ended, as recorded by the source sensor device or, possibly, a secondary source monitored by that device.

Timestamps for Resources

Timestamps are applied to the resources you see in the Navigator panel.

Timestamps for Resources	Context
Resource Created	This timestamp is applied by the Manager's JVM (Java Virtual Machine) when a resource is created.
Resource Modified	This timestamp is applied by the Manager's JVM (Java Virtual Machine) when a resource is changed.

Timestamp Variables

For date and time data fields, such as Detect Time, you can type an actual date value, such as 10/12/2016 8:54:00 AM, or can use special system variables such as:

- `$CurrentDateTime`: for the date and time the report is run; the system variable is replaced by the current date and time value.
- `$CurrentDate`: for the date the report is run; the system variable is replaced with the date value, truncating the time of the day to 0, when the report is scheduled or run.

You can also specify certain date operations with these system variables to add or subtract a number of specified days or hours. For example, you could type: `$CurrentDate - 7d` for seven days before the date the report is run, the condition evaluates to a date which is the current date minus seven days, or `$CurrentDateTime - 12h`, which evaluates to the current date time minus 12 hours. Do not create an operation that will result in a time stamp that is out of range

(the range is January 1,1970, at 00:00:01 through January 19, 2038 at 03:14:07, UTC), or you will get an error.

The time and date editing window you access through the **Detect Time** and **Detect Time Offset** fields of the ArcSight Console's Report Editor can accept month (uppercase **M**), minute (lowercase **m**), and current week (uppercase **W**) parameters.

Use spaces to separate these special system variables or parameters from other operators when including them in a condition statement.

See also ["Variables" on page 1064](#) and the subtopic on ["Timestamp Functions" on page 1077](#).

Inclusive Timestamps

The Detect Time timestamps reported for **correlated** events include the timestamps of the **base** events that initiated them. The timestamp is that of the most recent base event in the series of base events that caused the correlated event.

For example, an event's Detect Time field in the Event Inspector might now show 22 Sep 2017 18:18:24 PDT instead of 22 Sep 2017 16:10:29 PDT, with the difference being that the earlier timestamp represents the last base event rather than a later correlated event.

This refinement helps you interpret correlated events more readily, without the need to trace back through detailed rule chains.



Note: You can also inspect the Connector Time parameter to find out just when a rule triggered (the time that was recorded as the Detect Time in prior releases.)

Time Zone Correction

The correction of a local time zone is the number of hours of offset to apply in order to adjust local time to another clock (often UTC or GMT) to synchronize device-time queries, correlation, and filters.

Understanding Trends and Queries

A base trend is made up of one query. Trends can be used as the primary data source for a report. Or, a trend (based on one query) can be used as the data source to another query (a trend query) that further refines the initial query result. A collection of trend queries (queries that use trends as their data source) can provide focused views of a data set which can then be fed into a single report or multiple reports.

The system evaluates source data for trends based on event conditions (such as number of worm outbreaks, incident time-to-close, or number of cases closed) or common network elements (such as operating system, business role, or regulatory compliance relevance).

This provides a means of querying not just the current model of the network but to build reports on queries of historical data, scheduled queries, and snapshot trends. Using queries in trends allows you to evaluate, for example, trending statistics on vulnerabilities and incident metrics over time to determine whether your vulnerability posture or incident closing rate is getting better or worse.

The generic trend reporting and a set of specific reports show trends on current data. For example, you can evaluate trends by operating system, by role, by compliance requirement, time to close on cases, and number closed.

You can provide a trend on a selected period of time, and pull reports that generate aggregated data. Trends can include case metrics such as time to close, open and time open, number closed, which allows for trending reports on incidents.

For more information, see ["Query-Trend Relationships in Reporting" on page 468](#) in ["Building Trends" on page 467](#).

User Groups

User groups are named and organized collections of [Users](#). You can create groups based on departments, permission levels, work shifts, or whatever structure best supports your enterprise.

All users within a group inherit the group's permissions. If permissions are given to or taken from a group, all users within that group gain or lose those permissions. When users belong to more than one group, they receive permissions from all their groups. For example, if a user is in a group that has inspect permissions to all rules and is in another group with inspect permissions to all reports, the user will be able to inspect both rules and reports.

The following pre-defined groups help you manage your users:

- **Users:** Lists the current logged-in user and grants permission to inspect and edit their own information.
- **Shared:** Lists groups and users that the logged in user has permissions to.
- **All Users:** Lists all groups and users, only Administrators have permission to this group.

Groups created from the All Users group inherit permissions to only a few resources. You can either edit the group ACL to add or remove permissions or create groups beneath one of the pre-existing groups to inherit a pre-configured set of permissions.

- **Default User Groups:** Lists groups and users with default permissions to all resources. For more information on resources, see ["Editing Access Control Lists \(ACLs\)" on page 107](#).
- **Administrators:** Lists groups and users with full rights and access to manage all groups and users



Note: Do not delete the Administrators group. It grants administrative access. The Administrators group contains at least one user account. This user account is created during installation.

- **Live Rules Editors:** Lists groups and users with permissions to inspect and edit rules
- **Reports Editors:** Lists groups and user with permissions to inspect and edit reports
- **Unassigned:** Lists users who do not belong to a group

Users

Users are individuals who are assigned login names, passwords, and privileges to access and perform operations using the ArcSight Console or Command Center. For details on using the ArcSight Console for various tasks on dealing with users as an administrator, see ["Managing Resources" on page 711](#).

You manage users by storing user information, setting passwords, enabling or disabling login functionality, and organizing them into groups. When you create a new user account, a temporary password must be created for the user to login to the ArcSight Console. The user should change their password during their initial session. For more information on changing passwords, see ["Changing Your Password" on page 48](#). If you are an administrator, also refer to ["Resetting User Passwords" on page 104](#).

As an added security feature, user logins can be disabled. This feature may be used when the user is on an extended leave of absence, if the user ID and password have been compromised, or for any reason the user ID and password should not be used to access the ArcSight Console.

When users are deleted, they are removed from the Users resource tree but not from the database. The deleted user ID is stored in the database for future offline processing and user activity auditing. If the user belongs to more than one group, the user account is deleted from all groups automatically.

User Types

User accounts serve several purposes. To enable giving all users only the minimum set of privileges that are needed for them to fulfill their duties, user accounts have a "user type." The user type specifies, at a high level, which features a user may access. This mechanism is

complementary but does not replace permissions specified by access control lists (ACLs), which allow administrators to control access to resources such as assets, rules, and filters. User types are used primarily to control access to Manager services such as archiving and other management tools. See ["Managing Users and Groups" on page 99](#) and ["Managing Permissions" on page 107](#).

Most often, user types are used to limit the risk resulting from the fact that user name and password combinations are stored on disk for components that require unattended startup but have to authenticate to the Manager. For example, the Forwarding Connector needs to authenticate to the Manager in order to obtain events, but does not need access to any of the resource management functionality provided by archive and other management tools.

The currently supported user types are:

- **Normal User:** Has full privileges to use the Command Center and ArcSight Console, and all tools. Only apply this user type to accounts that actually need access to the ArcSight Manager.
- **Management Tool:** Has only the privileges needed to run certain management tools used in conjunction with network management products.
- **Forwarding Connector:** Has only the privileges needed by the Forwarding Connector.
- **Archive Utility:** Has only the privileges needed to run the archive utility. Access to specific resources is controlled through ACLs.
- **Connector Installer:** A specialized identity used only to add SmartConnectors to the system.
- **Web User:** Has privileges to use the Command Center only, not the ArcSight Console.

Unassigned users are those that do not belong to a group.

Variables

Variables are used to derive values from events, assets, and other resources (for example, a target IP address in an attack event, the MAC address or zone of a vulnerable asset, the timestamps on a user login session, entries in a hot list, and so forth).

You can use variables to create and tune [Active Channels](#), [Filters](#), [Reports](#), [Rules](#), [Field Sets](#), and [Data Monitors](#), or to expose more information, such as in report or grid view columns. The editors for these tools each include a **Variables** tab on which to add, edit, or remove variables.

Once created, variables appear in the [Common Conditions Editor \(CCE\)](#) as additional fields on the Filters or Conditions tabs; in Group By arguments for data monitors and rules; and in Select, Group By, and Order By fields for queries. In the Field Set Editor, variables are an additional category that appears once variables are defined.

Variables are especially useful for situational-awareness applications such as reporting on attacks by division, or for compliance monitoring as in reporting the number of compromise events directed at Sarbanes-Oxley related devices.

Asset-category variables are based on the relevant resource ID of the modeled network asset (device). Timestamp variables are based on the start, end, or receipt times recorded by SmartConnectors, Managers, or devices.

About Remote Variables

Variables using Group, List, and Category Model functions are evaluated and processed on the Manager, not directly on the ArcSight Console, and are referred to as *remote variables*.

These remote variables are evaluated only once on the ArcSight Console for any given event or resource. Therefore, the value of the variable on the ArcSight Console will not change if the underlying data is modified that would result in a different value for the variable. New events in events channels and resources in resource channels will evaluate the variable again, and you will see the updated value.

Because not all variables can be calculated on the ArcSight Console, there may be a delay in returning values from variables calculated remotely on the Manager.



Note: Variables do increase processing overhead and can affect report-generation performance. Consider the performance sensitivity of a report before adding variables.

About Functions

Functions configured for variables let you perform various operations on the derived values. To access event fields for use in **variable functions**, you either use the pick lists provided in the local or global variable dialogs or, in some cases, employ **velocity expressions** (templates) in statements. (See "[Velocity Templates](#)" on [page 1088](#) for an explanation of how to construct velocity expressions.)

When you click **Add** in a Variables tab, the Add Variable dialog box can present several fields, depending on the function to be used. All field values can be edited later except the choice of function. To change a Variable from one function to another, create a new Variable and delete the old Variable.

The Add Variable dialog includes the option to **preview** (or calculate) the results for some variable functions, given test values that you specify.



Note: Previewing the result of your variable definition

- The **Preview** (or Calculate) feature on the **Add Variable** dialog is supported for some but not all variable functions. For example, functions for list data types such as `GetSizeOfList` and `GetListElement` do not support the Preview feature.
- There is no way to specify a NULL value for **Preview** input to a Variable function. The Preview assumes that a blank field for an input is an empty String. Therefore, you cannot use Preview on the Variable dialog to test inputs for a parameter with NULL values.

See also "[Variable Definition Fields](#)" below.

Local and Global Variables

Variables you create in resources on the **Local Variables** tab of a resource editor are local to the resource for which you create them. For example, if you create a local variable in a query to get an active list value, that variable is available only to that query and not in other queries, rules, or filters. Queries themselves are available for use in trends, reports, and query viewers, but the local variables used to build them are not.

You can create global variables on field groups that are available across resources. The general information provided in this reference topic on variables, variable fields, and functions applies to both local and global variables. The main difference is that global variables are available across resources whereas local variables are not.

To create a local variable, click the **Variables** tab in a resource (for example, [Active Channels](#), [Filters](#), [Queries](#), [Rules](#)), name it, select a function, and provide arguments as needed.



Note: Ensure that local variable names are unique across resources. Local variables cannot share names.

To create a global variable, navigate to Field Sets, click the **Fields & Global Variables** tab, and use the Global Variables editor to select a function and parameters (as described in [Global Variables](#)).

Both local and global variables give you access to the same functions. Available functions are described in [Variable Definition Fields](#).

Variable Definition Fields

A variable has a name, a function associated with the variable, and one or more arguments.

Field	Description
Name	A meaningful name for the variable that is unique to the associated resource. Variable names must start with a letter, and can contain letters, numbers, underscores, and spaces. Trailing spaces at the end of a variable name will be removed. Special characters, other than those mentioned above, are not allowed.
Function	Functions are grouped into the following types: • "Alias Functions" below • "Arithmetic Functions" below • "Category Model Function" on page 1070 • "Condition Functions" on page 1072 • "Group Functions" on page 1073 • "IP Address Functions" on page 1074 • "List Functions" on page 1075 • "String Functions" on page 1076 • "Timestamp Functions" on page 1077 • "Type Conversion Functions" on page 1080 • "Value List Functions" on page 1084
Arguments	The contents of the Arguments section vary based on the Function selected. Functions require one, two, or three data fields as input arguments. The event data field list is filtered to show only fields of the required argument type. For example, the GetMonth function requires a single argument of type <code>timestamp</code>, so the list only shows timestamp-related fields: Agent Receipt Time, Device Custom Date 1, Device Custom Date 2, Device Receipt Time, End Time, Event Annotation Modification Time, and so on.

Alias Functions

Alias Function	Description
AliasField	Creates an alias (alternate name) for the specified field. 1. Provide the alias name you want to use. 2. Select a field from the drop-down list under Arguments.

Arithmetic Functions

The following table describes arithmetic functions. The binary arithmetic functions like Add, Subtract, and Multiply return a result type that is the higher resolution type of the two parameters. For example, `Add(Integer, Long)` returns `Long`. `Multiply(Integer, Double)` returns `Double`. The `Divide` function always returns `Double`.

Arithmetic Functions

Function	Description
Absolute	Returns the absolute value (its numerical value without regard to its sign) of the numeric argument. Arguments are integer, long integer, or double.
Add	Returns the result of adding the two numeric arguments together. Arguments are integer, long integer, or double types.
Ceil	Returns the smallest integer value that is not less than the numeric argument. Arguments are integer, long integer, or double.
Divide	Returns the result of dividing the first numeric argument by the second numeric argument. Arguments are integer, long integer, or double types; but the second argument must not evaluate to 0.
Floor	Returns the largest integer value that is not greater than the numeric argument. Arguments are integer, long integer, or double.

Arithmetic Functions, continued

Function	Description																			
Java Mathematical Expressions	Returns the result of the evaluation of the specified Java expression. Java Mathematical Expressions are for advanced users. ESM does not provide error checking and messaging for your JEP (Java expression parser) expressions. Refer to the Java math expressions parser web pages at http://www.singularsys.com/jep/ for more information on writing these expressions. Supported Expressions: ESM supports a subset of Java mathematical expressions parser (JEP), which are written like standard mathematical expressions. A JEP expression has three components: operator, function, and value. • Operator - Examples of operators are + - / * JEP operators are documented in the table on this Web site: http://www.singularsys.com/jep/doc/html/operators.html • Function - See http://www.singularsys.com/jep/doc/html/functions.html ESM supports the following functions, which are a subset of functions described in http://www.singularsys.com/jep/doc/html/functions.html. Supported Trigonometric Functions: <table><tr><td>sin(x)</td><td>cos(x)</td><td>asin(x)</td></tr><tr><td>acos(x)</td><td>atan(x)</td><td>atan2(y, x)</td></tr><tr><td>sinh(x)</td><td>cosh(x)</td><td>tanh(x)</td></tr><tr><td>asinh(x)</td><td>acosh(x)</td><td>atanh(x)</td></tr></table> Supported Logarithmic and Exponential Functions: <table><tr><td>ln(x)</td><td>log(x)</td><td>exp(x)</td></tr></table> Miscellaneous Supported Functions: <table><tr><td>abs(x)</td><td>rand()</td></tr><tr><td>mod(x,y)= x % y</td><td>sqrt(x)</td></tr></table> • Value - The values are either constants of numeric type or fields, which are referenced by the camelCase notation, such as bytesIn. For information on how to reference fields, refer to the "Script Alias" names in "Data Fields" on page 910. Example The expression "(bytesIn^2)/1000" squares the bytesIn value of an event and divides the result by 1000.	sin(x)	cos(x)	asin(x)	acos(x)	atan(x)	atan2(y, x)	sinh(x)	cosh(x)	tanh(x)	asinh(x)	acosh(x)	atanh(x)	ln(x)	log(x)	exp(x)	abs(x)	rand()	mod(x,y)= x % y	sqrt(x)
sin(x)	cos(x)	asin(x)																		
acos(x)	atan(x)	atan2(y, x)																		
sinh(x)	cosh(x)	tanh(x)																		
asinh(x)	acosh(x)	atanh(x)																		
ln(x)	log(x)	exp(x)																		
abs(x)	rand()																			
mod(x,y)= x % y	sqrt(x)																			

Arithmetic Functions, continued


Function	Description
	Notes: • All JEP functions return a Double. • Unlike velocity references, JEP expressions do not use "\$" in front of (ArcField) "Data Fields" on page 910. • Do not include mathematical operators or JME function names <i>in a variable name</i>. If you do, the JME parser interprets them as operators and returns unexpected results. Variable names that match JME function names such as sqrt cause similar problems. • Some expressions may not be valid and do not produce results. Do not use them in queries or active channels, and filters that use them cannot be used in queries or active channels. • JME variables are held only in memory and so, can be used only in Rules, Filters, and Data Monitors. They cannot be used in resources like Reports, which rely on persisted data. (There is a set of velocity references specifically for use in Reports. See "Velocity References for Reports" on page 1091 for more information.) • This function is held in memory, therefore you can only use it in Rules, Filters, and Data Monitors. You cannot use the function in resources like Queries and Reports, and other resources that rely on persisted data.
Multiply	Returns the product of multiplying the two numeric arguments together. The arguments may be integer, long integer, or double types.
Round	Returns the closest integer to the numeric argument. The argument must be a double.
RoundN	Takes two arguments, the double number to be rounded off and the number of decimal places from 0 to 5. Rounds off the input double number to the specified number of decimal places. For example: If the double is 1.23456789 and you want to make it readable in your reports, use the RoundN function to set the decimal places to 2. The double is rounded off to 1.23.
Subtract	Returns the result of subtracting the second numeric argument from the first numeric argument. The arguments may be integer, long integer, or double types.

Category Model Function

This function applies to Actors.

Category Model Function	Description
HasRelationship	Tests whether two actors, or an actor and a group, have the specified relationship based on a given Category Model. Category Model: Select an existing category model. Parent Field or Group: Select a field or single-value variable you want to use as the parent. Use the Field/Group drop-down to indicate whether the parent is a field (single attribute) or a group. Child Field or Group: Select the field or single-value variable you want to use as the child. Inherit All Related Actors: Select true to include all the actors in the selected group and its children. Select false to include only the actors in the selected group.

Condition Functions

Condition Functions	Description
ConditionalEvaluation	The ConditionalEvaluation function takes three arguments: - A filter defined as a Filter resource that acts as the conditional expression, - A value to return if the expression evaluates to True, and - A value to return if the expression evaluates to False.
CustomConditionalEvaluation	The CustomConditionalEvaluation function takes three arguments: - A local filter (not defined in the Filters resource) that acts as the conditional expression, - A value to return if the conditions evaluate to True, and - A value to return if the conditions evaluate to False. Enter the filter statement in the text field. Then create and edit local filter conditions through the Common Conditions Editor or Global Variables Selector: 
ReplaceNull	The ReplaceNull function takes two arguments: - A value to return from a test field, and - If the value is Null, a replacement value to return from another field. Both test and replacement fields must be of the same type. You can also set the replacement to a constant value. For example, for a String field, the replacement value can be something like Match found.

Group Functions


There are two general types of Group functions: FormatGroups and GetGroups.

The FormatGroups functions, FormatGroupsOfAssets and FormatGroupsOfNetworkZone, return a human-readable list of asset-category URIs unexclusively. This means that all matching and related categories are included. These variable functions mainly format and display asset category-groups. They are best used with the contents of fieldsets, reports, and data monitor fields. Avoid using the FormatGroups functions in conditions because result order cannot be assured for multiple-item groups; instead, use the GetGroups functions for ordering and consistency.

Group Functions	Description
FormatGroupsOfAsset	Formats the presentation of one or more matching asset category or asset group. The results are unordered. More Options on the Global Variable Editor's Parameters tab: If Anchor Point is <i>start</i> or <i>base</i>, the Offset is the number of elements to move from the anchor point when formatting the matched category URI. If Anchor Point is <i>end</i>, the Offset is the number of elements to move to the left. For example, if Asset1 is classified as A11 Assets/Location/USA/CA/Cupertino, the corresponding offset formats would be: USA = offset from Anchor Point by 0 CA = offset from Anchor Point by 1 Cupertino = offset from Anchor Point by 2 For MaxCount, specify the maximum number of categories to return. If the specified number is less than the number of matching categories, the returned categories are chosen at random. A Max Count of 0 means you want all categories returned.
FormatGroupsOfNetwork Zone	Formats one or more matching network zone group or asset category group. The results are unordered. See FormatGroupsOfAssets description for More Options.
GetGroupOfAsset	Returns a single Asset Category or Asset Category Group, given a Base Field and Base Group. If there is more than one matching category or group, a single URI is chosen at random. Related categories are not included. Output is optimized for correlation operations. Note: The GetGroups (plural Groups) functions return lists of asset categories, therefore their results cannot be used in inGroup conditions. This GetGroup (singular Group) function makes it possible to select one result at random, provided the variable is defined to produce a single result.

Group Functions	Description
GetGroupOfNetworkZone	Returns a single zone category. If multiple matches occur, a single URI is chosen at random. Related categories are not included. Note: The GetGroups (plural Groups) functions return lists of zones, therefore their results cannot be used in inGroup conditions. This GetGroup (singular Group) function makes it possible to select one result at random, provided the variable is defined to produce a single result.
GetGroupsOfAsset	Returns a list of Asset Categories or Asset Category Groups, given a Base Field and Base Group. In rule and data monitor aggregations this should produce multiple sets. In reports, this produces a coma-separated list of asset-category names. No related categories are excluded. Output is optimized for correlation operations. This function complements the FormatGroups functions. It simply shows XML representations of asset categories. Use this function in conditions and in Group By elements of rules or reports because its output is both well-ordered and consistent.
GetGroupsOfNetworkZone	Returns a list of Network Zone Groups or Asset Category Groups, given a Base Field and Base Group. This function complements the FormatGroups functions. It simply shows XML representations of group resources. Use this function in conditions and in Group By elements of rules or reports because its output is both well-ordered and consistent.

IP Address Functions



IP Address Function	Description
ParseIPAddress	Applies to IPv4 and IPv6 addresses. Takes two arguments: 1. The IP address 2. The byte position from 1 to 16. Each byte consists of two hex digits. • Positions 1 to 4 apply to IPv4 addresses. • Positions 1 to 16 apply to IPv6 addresses. Returns a value from 0 to 255. If you enter a byte position from 5 to 16 for an IPv4 address, the returned value is 0. IPv4 address examples: <pre>ParseIPAddress(192.0.2.27, 1) returns 192.</pre><pre>ParseIPAddress(192.0.2.27, 4) returns 27.</pre> IPv6 address examples: <pre>ParseIPAddress(2001:0db8:85a3:0042:1000:8a2e:0370:733b, 1) returns 32, the decimal format of the first two hex digits, 20 (one byte).</pre><pre>ParseIPAddress(2001:0db8:85a3:0042:1000:8a2e:0370:733b, 16) returns 59, the decimal format of the last two hex digits, 3b (one byte).</pre>

List Functions



Caution: Make sure you are getting values from case-sensitive lists only. Getting values from case-insensitive lists will negatively affect performance.

List Functions	Description
GetActiveListValue	Returns the value associated with a specific field of the specified active list. See "Active List Column Types and Subtypes" on page 490 for a list of supported fields.
GetSessionData	Returns the value associated with a specific field of the specified session list. Use this function for event and non-event schemas (any resource schema like actors, trends, cases, and so on), and specify the time at which the session is evaluated using either a time field, a constant time, or a dynamic time. See "Session List Column Types and Subtypes" on page 502 for a list of supported fields. See "Creating a Variable to Get Session List Data" on page 568 for an example of how this function is used.

String Functions

String Functions	Description
Concatenate	Returns the string result of joining the two string arguments. For example, <code>Concatenate("Arc", "Sight")</code> returns <code>ArcSight</code>. When a rule using this function fires and your string values have beginning or trailing spaces, the spaces are dropped, even if the preview during function definition or export to XML displays the space. If you want to enforce a space, use <code>ConcatenateThree</code> (described next). Note: This function is held in memory, therefore you can only use it in Rules, Filters, and Data Monitors. You cannot use the function in resources like Queries and Reports, and other resources that rely on persisted data.
ConcatenateThree	Returns the string result of joining the three string arguments. For example, <code>Concatenate("ArcSight", "Command", "Center")</code> returns <code>ArcSightCommandCenter</code>. Note: This function is held in memory, therefore you can only use it in Rules, Filters, and Data Monitors. You cannot use the function in resources like Queries and Reports, and other resources that rely on persisted data.
EvaluateVelocityTemplate	For advanced users with thorough understanding of velocity templates. Evaluates the velocity template argument and returns the result. This function is not available in a Query or Active Channel, and Filters that use this function cannot be used in a Query or Active Channel. For information on how to use Velocity Templates in ESM, see "Velocity Templates" on page 1088. Note: This function is held in memory, therefore you can only use it in Rules, Filters, and Data Monitors. You cannot use the function in resources like Queries and Reports, and other resources that rely on persisted data.
IndexOf	Returns the integer offset into the first string argument that is the location of the second string argument. For example, <code>IndexOf("Twas the night before Christmas", "night")</code> returns 9. If the second string argument is not found in the first string argument, <code>IndexOf</code> returns -1.
LastIndexOf	Returns the index (position) of the last (rightmost) occurrence of the second argument (the substring) within the first string argument (the source). If the substring is not found in the source, the function returns -1. The first position is index 0, as in the <code>indexOf</code> function. Examples: <code>lastIndexOf("abc/def/xyz", "/")</code> returns 7 <code>lastIndexOf("abc/def/xyz", "abc")</code> returns 0 <code>lastIndexOf("abc/def/xyz", "klm")</code> returns -1
LengthOf	Returns the number of characters in the string argument. For example, <code>LengthOf("Twas the night before Christmas")</code> returns 31. <code>LengthOf("")</code> returns 0.

String Functions	Description
Substring	Returns a portion of the first string argument, starting with the position specified in the second, numeric, argument and including the ending position as the sum of the number of characters and the starting position, specified in the third, numeric, argument. For example, Substring("Twas the night", 5, 8) returns "the".
ToLower	Returns the string argument converted to all lowercase. For example, ToLower("Inline Filter") returns "inline filter". Numbers and other non-alphabetic characters are not affected.
ToUpper	Returns the string argument converted to all uppercase. For example, ToUpper("Inline Filter") returns "INLINE FILTER". Numbers and other non-alphabetic characters are not affected.

Timestamp Functions

ESM applies timezones according to the component, shown below:

Time Zone	Description
Default Time Zone	The Manager time zone
Agent Time Zone	The time zone of the Connector which sent the event
Original Agent Time Zone	The time zone of the first Connector in a possible chain of connectors which sent the event
Device Time Zone	The time zone of the originally-reporting device
Final Device Time Zone	The time zone of the device which reported to the original Connector



Caution: Discrepancies in values returned by Timestamp functions

With certain resources, you might observe some discrepancy in values returned by Timestamp functions and End Time if ArcSight Manager and ArcSight Console are in different timezones. Following are the scenarios where the discrepancy occurs:

- For query viewers and data monitors, a Timestamp function (for example, GetDayOfWeek) gets the value from the Manager's timezone, and End Time gets the value from the Console's timezone.
- For reports, End Time values and values returned by TimeStamp functions are consistent with the Manager's timezone.
- For active channels, End Time values and values returned by TimeStamp functions are consistent with the Console's timezone.

TimeStamp Functions	Description
GetCurrentTime	Returns the current time in the format DD Mo YYYY hh:mm:ss TIMEZONE, for example 25 Jun 2016 14:05:18 PDT The returned time is based on the client time.
GetDayOfMonth	Returns an integer from 1 to 31 to represent the day of the month, based on the selected timestamp
GetDayOfWeek	Returns an integer from 0 to 6 (0 is Sunday) to represent the day of the week, based on the selected timestamp. The associated day of the week (for example "Sunday") is displayed on the ArcSight Console. You can test the value returned by this function using numeric operations like > , < , >= , <= , = . For example, for a variable called "day" that contains the value returned by the GetDayofWeek function, you can create an AND logical operator that checks for a weekday with these conditions: • day >= Monday • day <= Friday
GetDayOfYear	Returns an integer from 1 to 366 to represent the day of the year, based on the selected timestamp.
GetHour	Returns an integer from 0 to 23 to represent the hour of the day, based on the selected timestamp.
GetMinute	Returns an integer from 0 to 59 to represent the minute of the hour, based on the selected timestamp.
GetMonth	Returns an integer from 1 to 12 to represent the month of the year, based on the selected timestamp.
GetYear	Returns an integer for the year based on the selected timestamp and displays it as a 4-digit integer.
TimeDifference	Returns the result of subtracting the second timestamp argument from the first timestamp argument, in a human-readable format.
TimeDifferenceInDays	Returns the result of subtracting the second timestamp argument from the first timestamp argument, in days.
TimeDifferenceInHours	Returns the result of subtracting the second timestamp argument from the first timestamp argument, in hours.
TimeDifferenceInMinutes	Returns the result of subtracting the second timestamp argument from the first timestamp argument, in minutes.

TimeStamp Functions	Description
TimeDifferenceInSeconds	Returns the result of subtracting the second timestamp argument from the first timestamp argument, in seconds.
TimestampGranularity Note: This function is held in memory, therefore you can only use it in Rules, Filters, and Data Monitors. You cannot use the function in resources like Queries and Reports, and other resources that rely on persisted data.	Returns timestamp values at a granular level. This function is only available for in-memory operations like rules, data monitors, and channels; but not for reports, queries, and trends. Includes the following timestamp granularity options: get_year_only Returns a timestamp value of the first day of the year, first month of the year, and year; and zeroes out the hours, minutes, and seconds. For example, for a given timestamp of 4 Oct 2016 15:19:52 <Manager timezone>, the calculated value is 1 Jan 2016 00:00:00 <Manager timezone> get_year_month Returns a timestamp value of the first of the month, month, and year; and zeroes out the hours, minutes, and seconds. For example, for a given timestamp of 4 Oct 2016 15:19:52 <timezone>, the calculated value is 1 Oct 2016 00:00:00 <Manager timezone> get_year_month_day Returns a timestamp value of the date, month, and year only; and zeroes out the hours, minutes, and seconds. For example, for a given timestamp of 4 Oct 2016 15:19:52 <timezone>, the calculated value is 4 Oct 2016 00:00:00 <Manager timezone> get_year_month_day_hh Returns a timestamp value of the current date, month, year, and hours; and zeroes out the minutes and seconds. For example, for a given timestamp of 4 Oct 2016 15:19:52 <timezone>, the calculated value is 4 Oct 2016 15:00:00 <Manager timezone>

TimeStamp Functions	Description
	• get_year_month_day_hhmm Returns a timestamp value of the current date, month, year, hours, and minutes; and zeroes out the seconds. For example, for a given timestamp of 4 Oct 2016 15:19:52 <timezone>, the calculated value is 4 Oct 2016 15:19:00 <Manager timezone>
	• get_year_month_day_hhmmss Returns a timestamp value of the date, month, year, hours, minutes, and seconds. For example, for a given timestamp of 4 Oct 2016 15:19:52 <timezone>, the calculated value is 4 Oct 2016 15:19:52 <Manager timezone>
Note: You can test (click Calculate on the dialog for using this function in your variable) how each TimestampGranularity option calculates the value before you save the variable. The Manager's timezone is used in calculation.	

Type Conversion Functions

Click **Calculate** to view the conversion results in a Preview window.

ConvertAddressToString

Converts a given IPv4 or IPv6 address value to string. The input address must have a valid format: for example, the IPv4 address must have a valid number of octets, and the IPv6 address must have a valid number of bytes.

- IPv4 conversion examples as seen in the Preview window:

192.0.2.24 is converted to 192.0.2.24

::FFFF:192.0.2.24 (IPv4 embedded in IPv6) is converted to 192.0.2.24

Refer to [RFC 4291 Section 2.5.5](#) for information on IPv4-embedded addresses.

- IPv6 conversion examples as seen in the Preview window:

2001:0db8:85a3:0042:1000:8a2e:0370:7334 is converted to 2001:db8:85a3:42:1000:8a2e:370:7334 (leading zeroes are omitted)

2001:0DB8:0:0:0:0:0:0 is converted to 2001:db8:: (further simplified, and hexadecimal digits are converted to lowercase)

ESM returns null if you enter an invalid address format.

ConvertListToString

Takes as an argument the value of a multi-valued list entry and returns it as a comma-separated string (with each entry in the same format as displayed in a channel). This function works for both multi-valued active lists and session lists with overlapping entries.

For example, suppose you have a session list set up to show user names and IP addresses associated with login sessions. You would get user names from the session list via the `GetSessionData` variable. If there are three user names on the list (for example, darren, samantha, and endora), the `ConvertListToString` variable will return the three names (for example, `[darren, samantha, endora]`). You could do the same with IP addresses.

To use `ConvertStringToList`, first add a variable with the [GetSessionData](#) function. The nested fields that show up in the field selector (`<VariableNameFromSessionList>.<FieldName>`) can then be selected as arguments to this function.

For more about session lists, see ["Identity Correlation" on page 565](#) and ["List Authoring" on page 486](#).



Note: See also

- [ConvertStringToList](#)
- [List Functions](#)

ConvertNumberToString

Takes as an argument any number (integer, double, and so on) and returns it as a string.

ConvertResourceToReference

Takes a resource, for example, an asset, and converts it to a reference. Use this on an asset field you want, for example, Target Asset. This allows you to use a rule action to aggregate on that function, then use it to add an asset to an active list through the list's asset reference subtype.



Note: From an active channel containing the asset field, you cannot use this variable function to add the asset to an active list, even if the active list contains a resource reference field of subtype Asset. In other words, this variable function is not available for mapping assets from channels.

ConvertStringToDate

Converts a date and time pattern string to a timestamp format. Your string input formats can include the time in hours, seconds, and milliseconds; the AM/PM marker; and timezone.

Example input formats you can use:

month/day/shortyear or month/day/fullyear

You can optionally include the time in hours, seconds, and milliseconds; the AM/PM indicator; and timezone. For example:

```
mm/dd/yy hh:ss PM PST
```

You can specify the month by its name, for example March or the abbreviation Mar (case insensitive); or by its number, for example 03 or 3.

For a complete list of Java-specified formats supported by this conversion function, refer to <https://docs.oracle.com/javase/8/docs/api/java/text/SimpleDateFormat.html>.



Note: This function is held in memory, therefore you can only use it in Rules, Filters, and Data Monitors. You cannot use the function in resources like Queries and Reports, and other resources that rely on persisted data.

ConvertStringToDouble

Returns a double (floating point number) based on the selected string. For example, if a character string event field contained 3.19, ConvertStringToDouble would return a numeric value of 3.19.

ConvertStringToInteger

Returns a 4-byte integer based on the selected string.

ConvertStringToIPAddress

Takes an IPv4 or IPv6 address string and returns the corresponding binary IPv4 or IPv6 address value.

The IPv6 address input can be a full or compressed address string.

If your IPv6 string input contains blocks of 4-hex digits that are all zeroes, the function returns the compressed IPv6 address value. For example, if you enter

```
2001:db8:0000:0000:0000:0000:0000:0000
```

the function returns 2001:db8::

ConvertStringToList

Takes a comma-separated string and returns it as a multi-valued list. (See also [ConvertListToString](#).) You have the option to specify a separator string other than the default

comma, such as a pipe (|).

ConvertStringToLong

Takes an input string and returns a long (very large integer).

ConvertStringToMACAddress

Takes a MAC address string and returns the MAC address value. Your input must be a valid MAC address with hyphen separators:

```
00-00-5E-00-53-FF
```

ESM returns null if you enter an invalid MAC address.



Note: This function is held in memory, therefore you can only use it in Rules, Filters, and Data Monitors. You cannot use the function in resources like Queries and Reports, and other resources that rely on persisted data.

ConvertStringToResourceReference



Caution: This conversion function does not check the resource's validity. If the input string has the correct format, the type conversion takes place.

Takes an input string representing a resource ID or a resource URI and converts the string to a resource reference.

- Enter your Resource ID string in the following format (must match exactly):

```
<Resource ID="resourceIDvalue"/>
```

where *resourceIDvalue* is the unique 25-character value that conforms to ArcSight conventions for resource IDs. This ID is auto-generated and is shown on the resource's non-editable Resource ID attribute. For example

```
<Resource ID="QjZvvPPsAABCAEcWZ6-B1EQ==" />
```

Start the input string with a left angle bracket <, and end the string with a slash and right angle bracket />. Enclose the resource ID value with double quotes. Do not use a URI format for the resource ID.

- Enter your Resource URI string in the following format (must match exactly):

```
<Resource URI="/URI">
```

where */URI* is the URI to the resource. For example, the input can be

```
<Resource URI="/All Queries/ArcSight Administration/Connectors/System Health/Cache/Cache History by Connectors"/>
```

You can also combine Resource URI information with Resource ID in a single input string using the format:

```
<Resource ID="/MyURI" resourceIDvalue"/>
```

For example

```
<Resource URI="/All Queries/ArcSight Administration/Connectors/System  
Health/Cache/Cache History by Connectors" ID="QjZvvPPsAABCAEcWZ6-B1EQ==" />
```

Start the input string with a left angle bracket <, and end the string with a slash and right angle bracket >. Enclose the resource URI with double quotes, and enclose the resource ID with its own set of double quotes.



Note: This function is held in memory, therefore you can only use it in Rules, Filters, and Data Monitors. You cannot use the function in resources like Queries and Reports, and other resources that rely on persisted data.

Value List Functions

All functions in this category are held in memory, therefore you can only use them in Rules, Filters, and Data Monitors. You cannot use these functions in resources like Queries and Reports, and other resources that rely on persisted data.



Caution: Make sure you are getting values from case-sensitive lists only. Getting values from case-insensitive lists will negatively affect performance.

Value List Functions	Description
DistinctListValue	Takes a list and returns list elements, <i>excluding</i> null and duplicate values. The entries are enclosed in double quotes and separated by commas.
GetListElement	Takes two parameters, a list field and a list index (an integer), and returns the value from the specified <i>n</i> th index. The first list element is index 0. See "Using Functions: Examples with Lists" on the next page for additional information.
GetSizeOfList	Takes as an argument the value of a multi-valued list entry and returns the size of the list. For more about lists, see "List Authoring" on page 486. See "Using Functions: Examples with Lists" on the next page for additional information. Note: This function works for both multi-valued active and session lists with overlapping entries.

Value List Functions	Description
ListIntersection	Takes two lists and returns a single list containing values that are <i>common</i> to both lists, including null and duplicate values. Entries are enclosed in double quotes and separated by commas. Null values are represented as empty strings (""). If List1 and List2 each contain one null value and 1.0.0.1, this interested list is returned: <pre>"", "", "1.0.0.1", "1.0.0.1"</pre>
ListUnion	Takes two lists and returns a single list containing combined entries of both lists, including duplicates and null values. The entries are enclosed in double quotes and separated by commas. Null values are represented as empty strings ("").
NonNullListValues	Takes one list and returns a list of elements <i>except</i> null values. The entries are enclosed in double quotes and separated by commas.
SortListValues	Takes one list and returns a list of elements, <i>excluding</i> null values, sorted in ascending order. The entries are enclosed in double quotes and separated by commas.

Using Functions: Examples with Lists

Two examples are presented here:

- Getting login session data from a session list
- Extracting a list element from an active list

Getting Login Session Data from a Session List

Objective:

To get the number of login sessions maintained in a session list.

This scenario uses:

- Session list to be referenced by GetSessionData
- [GetSessionData](#) function specifying the session list from which to get values
- [GetSizeOfList](#) function that uses the GetSessionData variable as an argument
- [ConvertStringToList](#) function that uses the GetSessionData variable as an argument

We name the variable GetLoginsSessionData and use GetSessionData function. For this variable, specify the session list as the source of values. You can then select the nested fields that show up in the field selector:


```
(<VariableNameFromSessionList>.<FieldName>)
```

as the argument.

If there are three user names on the list (for example, darren, samantha, and endora), the `GetSizeOfList` function returns the number of names on the list (for example, [3]). You could do the same with the IP addresses.

After you have specified the field values to be taken from the session list, you can further use the `GetSizeOfList` and `ConvertStringToList` for additional methods to get your session data.

GetSessionData list function shows up as argument to ConvertListToString and GetSizeOfList Type Conversion functions.



Extracting a List Element from an Active List

Objective:

To extract the IP address from an active list containing expired audit events.

This scenario uses:

- Active list
- [ConvertStringToList](#)
- [GetListElement](#)
- [ConvertStringToIPAddress](#)

The scenario uses the value from `DeviceCustomString4`, where list elements are separated by a pipe (|):

```
desktop1.somecompany.com|mwhit|192.0.2.0|Antartica|ENG
```

In the string, the IP address is list element index **2**. To extract the IP address, create a chain of three variables as follows:

1. `parse_expired_entry = ConvertStringToList(DeviceCustomString4, "|")`
2. `get_ip_elem = GetListElement(parse_expired_entry, 2)`
3. `converted_ipa = ConvertStringToIpAddress(get_ip_elem)`

Variable Availability and Contexts

Not all variables are available in all contexts.

- These functions are only available for use with event schemas:
 - [ConditionalEvaluation](#)
 - [HasRelationship](#)
 - [AliasField](#)
- These functions are not available for use in SQL based operations:
 - [ConvertListToString](#)
 - [ConvertStringToList](#)
 - [GetSizeOfList](#)
 - [EvaluateVelocityTemplate](#)
 - [Java Mathematical Expressions](#)

Active Channels can evaluate [Group Functions](#), [Category Model Function](#), and [List Functions](#) only by sending a request to the Manager. Functions of these types are not evaluated on the ArcSight Console, unlike other variable functions. If you create active channels that use these function types, keep in mind that there will be a slight delay in an ArcSight Console channel display of these values.

See also "[Applying a Field Set to an Active Channel](#)" on page 236.

Variable Functions for In-Memory Operations

Functions listed below are used for *in-memory* operations only. This means you only use them on rules, filters, and data monitors. Such functions will not work on queries, reports, and active channels, which rely on persisted data.

- [Java Mathematical Expressions](#)
- [EvaluateVelocityTemplate](#)
- [Timestamp Functions](#)
- Some functions in [Type Conversion Functions](#)
- All functions in [Value List Functions](#)

Velocity Templates


ESM supports the use of *velocity templates* or scripts as defined by the Apache Velocity Project (<http://velocity.apache.org/>). Velocity templates are a means of specifying dynamic or variable inputs to, or outputs from, underlying Java code.

There are a number of places where a person familiar with Velocity templates can specify inputs using Velocity, instead of a literal value, to greatly enhance the results.



Caution: Velocity templates are for advanced users

- You must be experienced in using Velocity templates.
Because Velocity templates have such wide-ranging and intricate possibilities, mis-application or inappropriate application is entirely possible. OpenText cannot assume responsibility for adverse results caused by user-supplied Velocity templates.
- ESM does not provide error checking or error messaging for user-created velocity expressions. Refer to the Apache Velocity Project web page at <http://velocity.apache.org/> for more information on using velocity templates.
- Velocity template based variables are held only in memory and, therefore, can be used only in Rules, Filters, and Data Monitors. Velocity template based variables cannot be used in resources like Reports, which rely on persisted data. (There is a set of velocity references specifically for use in Reports. See "[Velocity References for Reports](#)" on [page 1091](#) for more information.)
- **Referencing Variables and Fields in Velocity Expressions.** Any variable that a velocity expression references must be local to the resource. You can refer to local variables and fields in a velocity expression.

If you have a global variable that you want to use in a velocity expression, use the +/-Global Variable button  on the [Common Conditions Editor \(CCE\)](#) to make it available in the resource. For more information, see "[Adding or Removing Global Variables Using the CCE](#)" on [page 901](#).

For more information on variables in general, see "[Variables](#)" on [page 1064](#) and "[Global Variables](#)" on [page 609](#).

Velocity Application Points

Velocity template support appears both in the user interfaces and in certain configuration files. The designated Velocity access points are described in the following table. Stated briefly, Velocity templates can be applied in most places where a literal string might be enhanced by a conditional or variable string. Common examples are formatting time expressions or condensing fine units into more meaningful groupings.

Velocity Template Usages

Application Point	Description
Rules Action Parameters	You can use Velocity templates in Add Action dialog boxes to create or edit fired-rule behavior. You get to these from the Actions tab or the Rules Editor. The Command and Parameters fields for Execute Command actions are Velocity candidates, as is the message-subject text in the Message field of Send Notification actions.
Custom Columns	Velocity templates are also applicable in the Cell Format and ToolTip Format panels of the Custom Columns Editor, which are described in "Customizing Columns" on page 250 .
SmartConnector Configuration	The URI strings in the Default Content tab of the Connector Editor can accept Velocity templates.
Case Audit Events	Audit events concerning cases can also be customized with Velocity templates, through properties files. In the case.default.properties or case.properties files (which overrides the former file), found at \$ARCSIGHT_HOME/config/audit, you can replace the expression in a key-value pair with a template variable or specify an additional field.
Notification Messages	In addition to using the Message field of Send Notification actions in the Add Action dialog box, you can also add Velocity templates to the destination-oriented notification configuration files located with the ArcSight Manager at \$ARCSIGHT_HOME/config/notification. This text controls message content (in contrast to the subject line).
Reports Text Fields	You can use a specific set of Velocity references for Report parameters when creating, editing, scheduling or running Reports and Focused Reports. Velocity references for Reports are covered in detail in "Velocity References for Reports" on page 1091 .

Using Velocity Expressions to Retrieve Values from Event Fields or Variables

Velocity expressions can be used to construct rule actions or velocity variables that need to access values in event fields or other variables. Rule actions can use velocity expressions in commands and notification messages. In these contexts, you need to write the velocity expression (there are no drop-down lists of fields provided, unlike in *rule conditions*). (See ["Managing Rule Actions" on page 532](#) and ["Rule Actions Best Practices" on page 538](#).)

You can construct most global variables and local variables simply by using the provided pick lists of event fields in the functions. However, the Arithmetic function [Java Mathematical Expressions](#) and the String function [EvaluateVelocityTemplate](#) are velocity variables that require you to write a velocity expression. (See ["Local and Global Variables" on page 1066](#).)

The syntax for constructing a velocity expression is the same, whether for rule actions or velocity variables.

Retrieving Values from Event Fields

To retrieve the value of an event field, use the field name in camel notation without any spaces, preceded by a dollar sign (\$):

```
$<fieldNameInCamelNotation>
```

For example, to retrieve the value of the Attacker Address field, use: `$attackerAddress`

For more about event fields, see ["Data Fields" on page 910](#).

Using Variables in a Velocity Expression

To retrieve the value of a variable, use the variable name preceded by a dollar sign (\$). If the variable name contains a dot, remove the dot and use camel case. If the variable name contains a space, use an underscore. See the following formats:

```
$<VariableName>  
$<variable_Name>
```

For example:

Variable display name	Velocity notation
Credit Card Number	<code>\$Credit_Card_Number</code>
dhcp.Hostname	<code>\$dhcpHostname</code>
Login User.Account Number	<code>\$Login_UserAccount_Number</code>

For more information, see ["Variables" on page 1064](#).

Using Velocity Expressions in Rule Actions

You can use velocity expressions in rule actions to retrieve the value of an event field or variable. These expressions can be used in commands or notification messages in rule actions.

For details syntax and guidance on constructing velocity expressions for use in rules, see ["Using Velocity Expressions to Retrieve Values from Event Fields or Variables" on the previous page](#).

Example of Rule Action that Uses Velocity Expressions to Retrieve Values

Following is an example of using both types of velocity expressions in a rule action to retrieve values from an event field (Attacker Address) and a variable (dhcp.Hostname):

1. In the Navigator panel, choose **Rules** from the drop-down menu.
2. Create or edit a rule.
3. Click the **Actions** tab.
4. Right-click a rule action and choose the **Send Notification** rule action.

The notification subject can be constructed as follows:

```
"Brute force login attempt from IP Address: $attackerAddress Hostname: $dhcpHostname"
```

5. Click **OK** or **Apply** to save the rule.

When the rule action is triggered, the notification message will replace the event field velocity expression "\$attackerAddress" with the value of the Attacker Address field, and the variable velocity expression "\$dhcpHostname" for the value of dhcp.Hostname.

Velocity References for Reports

The following Velocity references are available for use in [Reports](#) anywhere where text is used. These references pick up, contain, display, and print the given values. Generally, Velocity references in Reports are used for display and print purposes when creating, editing, scheduling or running Reports and Focused Reports. In some cases, they are used for more than that. For example in archived reports, \$Archive_Report_Folder and \$Archive_Report_Name determine the location where reports will be stored.



Note: The following table shows the complete set of applicable references for use with Reports. Other types of references (such as those discussed in the previous sections of this topic) do not apply to Reports. However, most of the details in [Velocity Template Usage Tips](#) also apply to Velocity Templates for reports.

Velocity References for Reports

Category	Reference	Description
Report	\$ReportName	Prints the name of the report, as specified in the Name field on the Attributes tab of the Report Editor.
	\$AccessDisclaimer	Prints a disclaimer statement regarding the user permissions with which the report was run. The disclaimer statement is a read-only string which is generated when report data has been filtered due to limited access privileges of the user Reports are generated only with data for which the current user has access privileges. Depending on user permissions for the user running a given report, access to some types of events or data may be curtailed. In such cases, the report is generated with all the information for which the user has access privileges. Events and data requiring higher-level access privileges are not included in the report. The access disclaimer statement is a standard explanation of the limitations of such a report.
	\$CurrentPageNumber	Prints the current page number of the report.
	\$TotalPageNumber	Prints the total number of pages in the report.
Time	\$CurrentDateTime	Prints the current date and time. (Same as \$Now) Example output: 12-06-2011-15:32:19. Tip: Formats for dates and times depend on your Console preference settings. To change the way dates and times are displayed throughout the Console, choose Edit > Preferences , then click the Date & Time button. For more information, see "Setting Date and Time Formats" on page 55 .
	\$CurrentDate	Prints the current date per your format preferences. Example output: 12-06-2011.
	\$CurrentMonth	Prints the current month. Example output: 12-2011.
	\$CurrentWeek	Prints the current week. Example output: 49-2011 (for December of 2011).
	\$Now	Prints the current date and time. (Same as \$CurrentDateTime) Example output: 12-06-2011-15:33:00.
	\$Today	Prints today's date. Example output: 12-06-2011-00:00:00.

Velocity References for Reports, continued

Category	Reference	Description
	\$CurrentDateTime- <Number>d	Prints the current date and time minus the number of days you specify. For example, if you ran the report on 12-06-2011 at 15:33:00 and specified the current date and time minus 1 day (\$CurrentDateTime-1d), this reference would output 12-05-2011-15:33:00. If, on the same day, you specified the current date and time minus 3 days (\$CurrentDateTime-3d), this reference would output 12-03-2011-15:33:00
Parameters	\$Report_Format	Prints the name of the report format that is configured as the default. Output formats are: • pdf - Adobe PDF file. • xls - Microsoft Excel file for tables and charts. (See "Setting Default and Custom Report Parameters" on page 415 for additional notes on the Report Format attribute, specifically on the XLS format.) • rtf - Rich-text format document • csv - Tabular data as a list of comma-separated values. • html - Web page displayed by the default web browser If the default output format for the report is set to html, then \$Report_Format reference simply will print the word html. See "Setting Default and Custom Report Parameters" on page 415 for information on how to set the default output formats for reports when creating reports. See also "Running a Report" on page 446 in for information about setting parameters at report runtime.
	\$Page_Size	Prints the page size of the report. Example output: Letter [8.5x11 in]
	\$Run_as_User	Prints the user name specified, if any, for the Run as User parameter in the report.
	\$Email_to	Prints the e-mail address specified, if any, for the Email to parameter in the report.
	\$Email_Format	Prints the e-mail format specified, if any, for the Email Format parameter in the report. For example, Send URL or Attach Report .
	\$Filter_by	Prints the filters used by the referenced query for this report.
	\$Archive_Report_Folder	Prints the folder location where the archived report is stored.
	\$Archive_Report_Name	Prints the name of the archived report.
	\$Archive_Report_Expiration_Time	Prints the expiration time for an archived report.

Velocity References for Reports, continued

Category	Reference	Description
	<code>\$<ComponentID>.Row_Limit</code>	Prints the row limit for the specified component. Tip: <code><ComponentID></code> refers to the data components or building blocks of a report. To view the components of a given report, right-click the report in the Navigator panel, choose Edit Report, and click the Data tab for the report. For example, if the report contains a component called Table, you can display related information by using the Velocity reference <code>\$Table.Row_Limit</code>, <code>\$Table.Time_Zone</code>, and so forth. Similarly, if the report, contains components called Chart1, Chart2, and Chart3; you can display related information on each of the charts by using references such as <code>Chart1.Time_Zone</code>, <code>Chart2.Start_Time</code>, and so forth.
	<code>\$<ComponentID>.Time_Zone</code>	Prints the time zone for the specified component. For example, <code>Table.Time_Zone</code> would output the time zone used for the data in a component called Table in your report. Example output: America/Los_Angeles
	<code>\$<ComponentID>.Start_Time</code>	Prints the start time for the specified component. For example, <code>Table.Start_Time</code> would output the start time used for the data in a component called Table in your report. (Start Time is a report parameter that can be configured on a per-component basis.) Example output: 12/05/2011 17:46:50.406-0800
	<code>\$<ComponentID>.End_Time</code>	Prints the end time for the specified component. For example, <code>Table.End_Time</code> would output the end time used for the data in a component called Table in your report. (End Time is a report parameter which can be configured on a per-component basis.) Example output: 12/05/2011 18:00:21.140-0800
	<code>\$<ComponentID>.<Parameter_name></code>	Prints the value of the specified component parameter.
	<code>\$Custom.<Parameter_name></code>	Prints the value a custom component parameter.

More Velocity Template Examples

You might use a Velocity template in a Zone URI field in an Connector Configuration Editor to specify a conditional target, as in:

```
#if( $deviceHostName.equals("foobar"))/All Customers/SuperCustomer#end
```

If you are setting up zones based on customers and you want to populate those values dynamically, you could use the following statement to populate fields based on host names, and so forth. For example, if you have one connector that collects events from devices monitoring different customers networks, you may want to set the customer name based on the device hostname.

```
device hostname = companyx.arcsight.com
```

The following template sets the customer name to **arcsight.com**:

```
CustomerURI=/All Customers/$deviceHostName.substring($deviceHostName.indexOf("."))
```

You can set the customer field from the SmartConnector as well, so events from a particular SmartConnector or device can be tagged as **customer xyz** (provided that Customer URI does exist on the Manager) and you can make ACLs limiting the customers' event privileges so they see only events tagged as **customer xyz**. If you have one SmartConnector that monitors devices reporting from multiple customers, you can dynamically set the customer name to be based on the device hostname. For example, if you have a customer named **arcsight** and the device hostname is **device1.arcsight.com**, the following template returns **arcsight** as the customer name:

```
CustomerURI=/All Customers/$deviceHostName.substring($deviceHostName.indexOf("."),$deviceHostName.lastIndexOf(".")).substring(1)
```

The result would be the URI: /All Customers/arcsight

For a case audit event in case.default.properties, a template could consist of:

```
deviceCustomString3=$history
```

Velocity Template Usage Tips

- **Use with strings and numeric values only.**

Velocity templates apply **only** to fields that contain string or numeric values.

- **Use with dynamic parameters and ArcSight variables.**

You can use all of the dynamic time parameters you see in the Active Channel Editor and elsewhere, such as \$Now and \$CurrentDateTime. The same is true for time elements, including s (second), m (minute), d (date), M (month), w (week), and y (year). To use any event data field as a variable, express its displayed name as a one-word, camel case string prefixed with a dollar sign. For example, "Source Address" would be \$sourceAddress. For details about using variables in a velocity expression, see ["Using Variables in a Velocity Expression" on page 1090](#).

- **Regular expressions are not supported.**

Use of regular expressions is not tested or supported.

- **Test using active channel custom field.**

You can conveniently test Velocity templates by trying them first in a customField of an active channel.

Views

"Views" is a collective term for all the different options you have for seeing raw and processed [Events](#) information in the ArcSight Console's Viewer panel.

The Console's Viewer panel can display event information in several formats and is readily customizable. Views may be customized to best reflect an enterprise and can be organized in a hierarchical structure with drill-down functionality. There is a list of chart-format views in addition to grids, maps, and dashboards.

See also ["Viewing" on page 66](#) and ["Monitoring Active Channels" on page 224](#).

View Types

Each view type represented by a tab at the **top** of the Viewer panel serves as a container for all individual instances of that type of view. For example, all data monitors opened in a dashboard remain part of it, and also inherit any visual choices you make for that view. Using the **View Layout** icon at the lower-right corner of the Viewer panel you can choose to tile or tab the individual views. When you tab the views, you select them using the tabs at the **bottom** of the panel.

The screenshot shows the ArcSight Console Viewer panel with the following sections:

- Latest Events By Priority**: A tab at the top left.
- Events By Priority**: A table with columns: Priority, Total (Total Legends 0).
- Latest Severe Threat Events**: A table with columns: Name, Categ..., Categ..., Categ..., Attac..., Targe..., Targe..., Priority.
- Latest Elevated Threat Events**: A table with columns: Name, Cate..., Cate..., Cate..., Attac..., Targe..., Targe..., Priority. It contains several rows of data with priority levels 5 and 7.
- Latest High Threat Events**: A table with columns: Name, Cate..., Cate..., Cate..., Attac..., Targe..., Targe..., Priority. It contains several rows of data with priority levels 7 and 8.
- Latest Guarded Threat Events**: A table with columns: Name, Cate..., Cate..., Cate..., Attac..., Targe..., Targe..., Priority. It contains several rows of data with priority level 3.
- Latest Low Threat Events**: A table with columns: Name, Categ..., Categ..., Categ..., Attac..., Targe..., Targe..., Priority.

At the bottom, it says "Data last refreshed: 7/21 13:02:33".

With views you have the flexibility to monitor an enterprise from various perspectives. Views can be customized to best capture and reflect an enterprise's network infrastructure and can also be organized in a hierarchical structure with drill-down functionality. Views can vary in scope and scale, from broad to detailed, depending on how the enterprise is monitored and organized.

The ArcSight Console provides different views in which you can display event data in the Viewer panel. You can select which views to display by selecting options from the Views menu.

Dashboards

Dashboards provide a more customized view of data, letting you create individual "instrument panels," each of which can display results based on different event data and filter conditions, and in different formats.

From the Viewer panel, you can change the view type or format of individual tabs from grid to line chart, bar chart, pie chart, or graphic. In addition, you can **float** the display of individual sub-view tabs, dashboards, and individual data monitors into separate windows to expand or resize individual displays.

While chart views display a summary of events, grid views display each event. Grid views display events organized in rows and columns. As new events occur, they are inserted at the top of the grid as a new row. Rows contain events while columns contain data fields.

Other Views

The Console automatically shows HTML information such as reports, references pages, and results for the Web Search tool in your default Web browser.

The Viewer panel is where you use the Find Resource query editor and result details. (See also ["Finding Resources" on page 704.](#))

Vulnerabilities

A vulnerability is a hardware, firmware, or software state that leaves an automated information system (AIS) open for potential exploitation. It could be due to anything, including circumstance, configuration, design, or implementation. A vulnerability can also be described as a weakness in automated system security procedures, administrative controls, physical layout, internal controls, and so forth, that could be exploited by a threat to gain unauthorized access to information or disrupt critical processing.

Vulnerabilities are discovered using scanners and their associated SmartConnectors. The Manager imports the output from vulnerability scanners, recording them as items in the Vulnerabilities resource tree, in the Assets section of the Navigator panel. Vulnerabilities are mapped to their associated devices. Vulnerabilities describe asset threats and exposures and provide more information with a link to Knowledge Base articles or notes.

Vulnerability Groups

Vulnerability groups are created to store similar groups of vulnerabilities in a single location. Groups can be created within groups to meet enterprise needs. When a group is created within a group, the new group inherits the existing group's permissions. If a group is deleted, the vulnerabilities within that group are also deleted. The following groups are provided:

- **Shared:** vulnerabilities to which logged-in users have permission.
- **Unassigned:** vulnerabilities that are not assigned to a group.

If you have Administrator access you will have another group named All Vulnerabilities that contains all vulnerability groups and vulnerabilities.

Standardized Vulnerability Tracking

In the Vulnerabilities tab of the Assets resource tree, there is a branch for using the MITRE Corporation's CVE (Common Vulnerabilities and Exposures) standardized vulnerability naming and reference system.

CVE is a list (dictionary) of standardized names for vulnerabilities and other information security exposures. CVE seeks to standardize the names for all publicly known vulnerabilities and security exposures.

You can map CVE as one of its vulnerability reference authorities, within its Navigator panel resource tree. This information can serve, for example, to determine the significance of IDS events. The goal of CVE is to provide a common naming scheme, shared by vulnerability scanners and other security devices to link real-time events to asset vulnerabilities.

You can search its CVE-related Navigator panel resources by CVE name, and to include CVE names in its ArcSight Console or report output.

The requirements for CVE compatibility are fulfilled by the capacity to analyze event streams utilizing CVE names, generate reports for CVE-related vulnerabilities, map events to asset vulnerabilities, and the existence of documentation for CVE-related functionality.

Web Browsers

You can launch HTML-based displays in an external Web browser from the ArcSight Console.



Note: Refer to the Support Matrix applicable to your ESM version for an official list of supported Web browsers.

Browser Preferences for HTML Displays

The ArcSight Console offers a general preference option for HTML display of various information in your preferred external Web browser.

The way you set these browser preferences determines display of reports, knowledge base articles, graphs and charts, and so forth. (For information on the general setting for HTML viewing preferences, see the table on [Setting Default Editors and Viewers](#) for information on preferred Web browsers.)

Browser Preference Overrides for Specific Features

Additionally, you can set your viewer preference for HTML displays specifically for certain features, and override the general preference setting for these specific displays. Some examples are:

- Integration command configurations. HTML display preferences for integrated command results are set as attributes on the command configuration. See "[Configurations Attributes](#)" on page 677 for more information.



Note: For ESM 6.5c and later, use the ArcSight Command Center to run Logger searches. Refer to the [ArcSight Command Center Users Guide](#) for information on running searches.

- Online Help. You can set a preference specific to the Online Help for display in an external Web browser.

Publication Status

Released: November, 2025

Send Documentation Feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this computer, click the link above and an email window opens with the following information in the subject line:

Feedback on ArcSight Console User's Guide (ESM 7.9)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to Documentation-Feedback@microfocus.com.

We appreciate your feedback!