

Administrator's Guide

ArcSight Logger 5.5

June 10, 2014



Copyright © 2014 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Follow this link to see a complete statement of copyrights and acknowledgements:

<http://www.hpenterprisesecurity.com/copyright>

The network information used in the examples in this document (including IP addresses and hostnames) is for illustration purposes only.

HP ArcSight products are highly flexible and function as you configure them. The accessibility, integrity, and confidentiality of your data is your responsibility. Implement a comprehensive security strategy and follow good security practices.

This document is confidential.

Contact Information

Phone	A list of phone numbers is available on the HP ArcSight Technical Support page: http://www8.hp.com/us/en/software-solutions/software.html?compURI=1345981#.URitMaVwpWI
Support Web Site	http://support.openview.hp.com
Protect 724 Community	https://protect724.hp.com

Revision History

Date	Product Version	Description
06/10/2014	5.5	Update for Logger 5.5
02/07/2014	5.5	5.5 GA release
02/12/2013	5.3 SP1	5.3 SP1 release
08/22/2012	5.3	5.3 GA release
12/05/2011	5.2	5.2 GA release
05/14/2011	5.1	5.1 GA release
11/08/2010	5.0 Patch 2	A special Patch 2 update for this guide to update remote access to Logger information.
09/16/2010	5.0	5.0 GA release
07/06/2010	4.5	4.5 GA release
01/22/2010	4.0 SP1	4.0 SP1 release
11/03/2009	4.0	4.0 GA release
07/22/2009	3.0 SP1	3.0 SP1 release
01/07/2009	3.0 Patch 1	3.0 Patch 1

Contents

Chapter 1: Overview	21
Introduction	21
Logger Features	23
Storage Configuration	23
Receiver Configuration	23
Analyzing Events	24
Grouping Events	25
Exporting Events	25
Forwarder Configuration	26
User Management	26
Other Setup and Maintenance	26
Deployment Scenarios	27
Centralized Management	28
What's New in Logger 5.5	28
Chapter 2: Installation and Initialization	29
Deployment Planning	29
Storage Strategy	29
Retention Policy	30
Initial Configuration	31
SAN	31
Storage Volume	31
Storage Groups	32
Indexed Fields and Full-text Indexing	32
Receivers	32
Initializing a Logger Appliance	33
Acquire a License for the Logger Appliance	34
Log In and Accept the License Agreement	34
Initialize the Logger Appliance	35
Set Up the Logger Appliance for Remote Access	39
Installing a Software Logger	39
Supported Platforms	39
Downloading the Logger Software	39
How Licensing Works in Software Logger	40

Acquiring a License for a Software Logger	41
Prerequisites for Installation	42
Installation Modes	42
Installation Steps	43
Using GUI Mode to Install Software Logger	43
Using the Console Mode to Install Software Logger	45
Using the Silent Mode to Install Software Logger	46
Starting and Stopping the Software Logger	48
Uninstalling Software Logger	49
Connecting to Logger for the first time	49
Configuring Logger	50
Receivers	50
Enabling the Pre-configured Receivers	51
Devices	51
Device Groups	52
Storage Rules	52
Using SmartConnectors to Collect Events	52
SmartMessage	53
Downloading SmartConnectors	53
Configuring a SmartConnector to Send Events to Logger	53
Configuring SmartConnectors to Send Events to Both Logger and an ArcSight Manager	54
Configuring SmartConnectors for Failover Destinations	54
Sending Events from ArcSight ESM to Logger	55
Chapter 3: User Interface and Dashboards	57
Connecting to the Logger User Interface	57
Logging In	57
Navigating the User Interface	59
Help	60
Options	60
About	61
Logout	61
Summary	61
Dashboards	64
Creating and Managing Dashboards	66
Adding and Managing Panels in a Dashboard	68
The Default Monitor Dashboard	70
Platform	72
Network	72
Logger	73
Receivers	73
Forwarders	73
Storage	74

Chapter 4: Searching and Analyzing Events	75
The Need to Search Events	75
The Process of Searching Events	76
Simple Query Example	76
Query Example Using a Chart	77
Elements of a Search Query	78
Query Expressions	78
Indexed Search	78
Search Operators	86
Time Range	86
Fieldsets	88
Creating Custom Fieldsets	89
Constraints	92
Syntax Reference for Query Expressions	93
Using the Advanced Search Builder Tool	96
Accessing the Advanced Search Builder	97
Nested Conditions	99
Alternate Views for Query Building in Search Builder	100
Search Analyzer	100
Performance Optimizations for Indexed Fields in Search Queries	101
Regex Helper Tool	101
Search Helper	103
Autocomplete Search	104
Search History	105
Search Operator History	105
Examples	105
Usage	105
Suggested Next Operators	106
Help	106
Searching for Events on Logger	106
Advanced Search Options	107
Searching Peer Loggers (Distributed Search)	108
Tuning Search Performance	108
Understanding the Search Results Display	109
User-defined Fields in Search Results	111
Viewing Search Results Using Fieldsets	112
Using the Histogram	112
Multi-line Data Display	113
Auto Updating Search Results	114
Chart Drill Down	114
Understanding Field Summary	116
Refining and Charting a Search from Field Summary	118
Exporting Search Results	119

Scheduling an Export Operation	121
Indexing	121
How Indexing Works	121
Full-text Indexing (Keyword Indexing)	122
Field-based Indexing	122
Guidelines for Field-based Indexing	124
Enabling Indexing	124
Adding Fields to Field-based Index	125
Super Indexing	125
Saving Queries (Saved Filters and Searches)	126
Saving a Query	126
Using a Saved Filter or a Saved Search	128
System Filters/Predefined Filters	128
Using a System Filter	132
Alerts	132
Viewing Alerts	132
Receiving Alerts for Events	132
Base Event Fields	133
Go, Export, and Auto Update Options	133
Live Event Viewer	133
Chapter 5: Reporting	137
The Reports Home Page	137
Explorers	139
Category Explorer	140
Report Explorer	141
Query Explorer	142
Parameter Explorer	143
Favorites Explorer	144
List of Buttons in the Explorers	144
Categories	147
System Defined Categories	147
Anti-Virus	147
Configuration Monitoring	148
Cross Device Reports	148
Database Reports	148
Default Reports	148
Firewall Reports	148
Identity Management Reports	148
IDS-IPS Reports	148
Intrusion Monitoring Reports	148
Network Reports	148
Network Monitoring Reports	149

Operating System Reports	149
SANS Top 5 Reports	149
VPN Reports	150
Solution Reports	150
Adding a New Category	151
Deleting an Existing Category	152
Placing a System Defined Query or Parameter into a Category	153
Dashboards	153
Viewing the Dashboard	154
Designing Dashboards	154
What Items Can a Dashboard Include?	155
Creating a New Dashboard	155
Dashboard Buttons	156
Viewing an Existing Dashboard in a Tab in the Dashboard Viewer	156
Removing an Existing Tab from the Dashboard Viewer	157
Deleting a Dashboard	157
Editing an Existing Dashboard	158
Selecting a Default Dashboard View for the Reports Home Page	158
Widgets	159
The Widget Designer	159
Creating a New Widget	159
Deleting a Widget	161
Editing an Existing Widget	161
Placing Widgets in a Dashboard	162
Moving an Existing Widget within a Dashboard	162
Using Dashboards Created in Pre-5.2 Logger	162
Viewing a Classic Dashboard	162
Designing Classic Dashboards	162
What Items Can a Dashboard Include?	163
Creating a New Classic Dashboard	163
Placing Items onto the Existing Dashboard	163
Creating Widgets	164
Placing Dashboard Items on the Layout	164
Swapping Items on Widgets	170
Setting Pre-5.0 Dashboard Preferences	171
Working with Available Dashboards	171
Selecting a Dashboard View	171
Running, Viewing, and Publishing Reports	172
Best Practices	172
Finding Reports	172
Viewing Recently Run Reports	173
Task Options on Available Reports	173
Running and Viewing Reports	174

About the Pagination of Reports	174
Quick Run with Default Options / Run In Background Report Parameters	175
Run Report Parameters	178
Report File Formats	179
Publishing Reports	180
Report Delivery Options	181
Refreshing a Report	182
E-mailing a Report	182
Exporting and Saving a Report	184
Viewing the Output of a Published Report	184
Designing Reports	185
Opening the Report Designer	185
Creating New Reports	186
Quick Start: Base a New Report on an Existing One	186
Designing New Reports	187
Select Filter Criteria	189
Select Grouping	191
Select Totals	192
Sort Order	192
Highlighting	193
Create Matrix	194
Create Chart	195
Editing a Report	197
Private Reports	197
Adhoc Report Designer	197
Setting Access Rights on Reports	198
Determining What Access Rights to Give a Group or User	199
Setting up Queries	201
How Search and Report Queries Differ	202
Overview of Query Design Elements	202
Creating a Copy of an Existing Query	202
Designing a New SQL Query	203
Modifying a Query Object	215
Deleting a Query Object	215
Defining SQL in the Editor	215
Working with Parameters	223
Creating New Parameters	224
Modifying a Parameter	229
Deleting a Parameter	229
Configuring Parameter Value Groups	229
Applying Report Template Styles	232
Defining a New Template	232
Scheduling Reports	233

Viewing and Editing Scheduled Reports	233
Scheduling a Report	234
Add Report Job Settings	236
Deploying a Report Bundle	236
Report Server Administration	238
Timeouts when Running Reports	238
Report Configuration	238
Using Report Category Filters	239
Backup and Restore of Report Content	240
iPackager	240
The iPackager Page	240
Buttons Available from the iPackager	241
Importing References from the Report Server	242
Modifying Properties for Imported Objects	243
Category Properties	243
Report Properties	243
Query Properties	244
Parameter Properties	244
Template Properties	245
Opening a .conf File	245
Deleting an Item from the .conf File	245
Clearing the Contents in a .conf File	245
Building the CAB	245
Deploying a CAB file in Logger	246
Chapter 6: Configuration	247
Devices	247
Devices	247
Device Groups	249
Event Archives	250
Guidelines for Archiving Events	252
Archiving Events	253
Scheduled Event Archive	254
Archive Storage Settings	255
Loading and Unloading Archives	256
Storage	256
Storage Groups	257
Storage Rules	259
Storage Volume	260
Event Input	260
Receivers	260
File Based Receivers	261
Working with Receivers	263

Source Types	274
Working with Source Types	275
Parsers	278
Using Parsers with Source Types	279
Using the Parse Command	279
Working with Parsers	280
Example: Creating an Extract Parser	282
Event Output	284
Forwarders	285
ESM Destinations	291
Certificates	294
Forwarding Log File Events to ESM	295
Alerts	296
Alert Triggers and Notifications	297
When are Alert events triggered?	297
Receiving Alert Notifications	298
Sending Notifications to E-mail Destinations	298
Sending Notifications to Syslog and SNMP Destinations	298
Configuring and Managing Real Time Alerts	299
Creating a Real Time Alert	299
Creating and Managing Saved Search Alerts	302
Creating a Saved Search Alert	303
Sending Notifications to SNMP Destinations	307
Sending Notifications to Syslog Destinations	308
Sending Notifications to ESM Destinations	309
Scheduled Tasks	309
Scheduled Tasks	310
Currently Running Tasks	310
Finished Tasks	311
Filters	312
Filters Tab	312
Search Group Filters Tab	314
Saved Searches	315
Saved Searches	315
Scheduled Saved Search	316
Saved Search Files	320
Search	320
Adding Search Indexes	320
Tuning Advanced Search Options	321
Viewing and Deleting Field Sets	323
Viewing Default Fields	324
Viewing Custom Fields	324
Running Search Tasks	325

Ending Currently Running Tasks	325
View and Add Parsers for Specific Log Types	326
Peer Loggers	326
Guidelines	326
Authorizing Peers	330
Configuration Backup and Restore	330
Running a Configuration Backup (Ad-hoc or Scheduled)	331
Restoring from a Configuration Backup	332
Editing Configuration Backup Settings	333
System Maintenance	333
Entering Maintenance Mode	334
Exiting Maintenance Mode	334
Checking Status of a Maintenance Operation	334
Database Defragmentation	335
Guidelines for Database Defragmentation	335
Defragmenting a Logger	336
Global Summary Persistence Defragmentation	340
Guidelines for Global Summary Persistence Defragmentation	340
Storage Volume Size Increase	342
About Increasing Storage Volume Size on a SAN Logger	342
Adding Storage Groups	343
Adding or Importing Schema Fields	345
Importing Schema Fields from Peers	347
License Information	351
Data Volume Restrictions	351
Retrieve Logs	352
Content Management	353
Importing Content	353
Importing Guidelines	354
Exporting Content	354
Exporting Guidelines	355
Chapter 7: System Admin - Logger Appliance	357
System	358
System Locale	358
System Reboot	358
Network	358
System DNS	359
Hosts	359
NICs	359
Static Routes	361
Time/NTP	361
SMTP	363

License & Update	364
Updating the Appliance	364
Updating the License File	364
Process Status	364
SNMP	365
Sending System Health Events as SNMP Traps	366
Polling System Health Information Using SNMP	366
Viewing Polled Information	367
SSH Access to the Appliance	370
Enabling or Disabling SSH Access	370
Connecting to Your Appliance Using SSH	370
Logs	371
Audit Logs	371
Audit Forwarding	371
Storage	372
Remote File Systems	372
Managing a Remote File System	373
SAN	375
Restoring a SAN	377
Creating Multiple Paths to a LUN	378
RAID Controller/Hard Disk SMART Data	379
Security	380
SSL Server Certificate	381
Generating a Self-Signed Certificate	381
Generating a Certificate Signing Request (CSR)	383
Importing a Certificate	385
SSL Client Authentication	386
Configuring Logger to Support SSL Client Authentication	386
Uploading Trusted Certificates	387
Uploading a Certificate Revocation List	387
Enabling Client Certificate Authentication	388
FIPS 140-2	388
Users/Groups	392
Authentication	392
Sessions	392
Local Password	393
Users Exempted From Password Expiration	395
Forgot Password	396
External Authentication	396
Login Banner	401
User Management	402
Users	402
Groups	404

Change Password	407
Other System Administration Information	407
Monitoring System Health	407
System Health Events	408
Using the Command Line Interface	410
Chapter 8: System Admin - Software Logger	413
System	413
System Locale	413
SMTP	414
License & Update	414
Updating the License File	414
Process Status	415
System Settings	416
Logs	416
Audit Logs	416
Audit Forwarding	417
Security	417
SSL Server Certificate	417
Generating a Self-Signed Certificate	418
Generating a Certificate Signing Request (CSR)	419
Importing a Certificate	421
SSL Client Authentication	422
Configuring Logger to Support SSL Client Authentication	422
Uploading Trusted Certificates	423
Uploading a Certificate Revocation List	424
Enabling Client Certificate Authentication	424
FIPS 140-2	424
Users/Groups	428
Authentication	428
Sessions	428
Local Password	429
Users Exempted From Password Expiration	431
Forgot Password	432
External Authentication	432
Login Banner	437
User Management	438
Users	438
Groups	440
Change Password	442
Other System Administration Information	443
Monitoring System Health	443
System Health Events	443

Chapter 9: Managing Connectors	445
Connector Overview	445
Navigating the Manage Connectors Tab	446
Locations	447
Viewing All Locations	448
Viewing Hosts, Containers, and Connectors in a Location	448
Adding a Location	448
Exporting and Importing Remote Management Configuration	449
Adding Locations and Hosts from a File	450
Editing a Location	451
Deleting a Location	451
Adding Hosts to a Location	451
Hosts	451
Viewing All Hosts	452
Viewing Containers and Connectors in a Host	452
Adding a Host	452
Scanning a Host	454
Deleting a Host	456
Moving a Host to a Different Location	456
Editing a Host	456
Upgrading a Host Remotely	457
Adding a Container to a Host	457
Containers	457
Viewing All Containers	458
Viewing Connectors in a Container	458
Adding a Container	459
Adding a Connector to a Container	459
Editing a Container	459
Deleting a Container	459
Updating Container Properties	460
Changing Container Credentials	461
Enabling and Disabling FIPS on a Container	461
Managing Certificates on a Container	462
Adding CA Certificates on a Container	463
Removing CA Certificates from a Container	464
Adding a CA Certs File on a Container	465
Enabling or Disabling a Demo Certificate on a Container	466
Adding Multiple Destination Certificates to a Container	467
Viewing Certificates on a Container	468
Resolving Invalid Certificate Errors	469
Running a Command on a Container	470
Upgrading a Container to a Specific Connector Version	470
Viewing Container Logs	471

Deleting Container Logs	472
Running Logfu on a Container	472
Running Diagnostics on a Container	473
Connectors	474
Viewing all Connectors	474
Adding a Connector	475
Editing Connector Parameters	478
Updating Simple Parameters for a Specific Connector	478
Updating Table Parameters for a Specific Connector	479
Updating Simple and Table Parameters for Multiple Connectors	480
Managing Destinations	481
Adding a Primary Destination to a Specific Connector	481
Adding a Failover Destination to a Specific Connector	483
Adding a Primary or Failover Destination to Multiple Connectors	484
Removing Destinations	485
Re-Registering Destinations	486
Editing Destination Parameters	487
Editing Destination Runtime Parameters	489
Managing Alternate Configurations	490
Sending a Command to a Destination	492
Removing a Connector	493
Sending a Command to a Connector	494
Running Logfu on a Connector	494
Changing the Network Interface Address for Events	495
Developing FlexConnectors	495
Editing FlexConnectors	498
Sharing Connectors (ArcExchange)	499
Packaging and Uploading Connectors	499
Downloading Connectors	502
Configuration Suggestions for Connector Types	504
Deploying FlexConnectors	504
Configuring the Check Point OPSEC NG Connector	505
Adding the MS SQL Server JDBC Driver	507
Adding the MySQL JDBC Driver	508
Chapter 10: Managing Repositories	509
Overview	509
Logs Repository	511
Uploading a File to the Logs Repository	511
CA Certs Repository	511
Uploading CA Certificates to the Repository	512
Removing CA Certificates from the Repository	513
UpgradeAUP Repository	513

About the AUP Upgrade Process	514
Uploading an AUP Upgrade File to the Repository	514
Removing a Connector Upgrade from the Repository	514
Content AUP Repository	515
Applying a New Content AUP	515
Applying an Older Content AUP	516
Remote Management AUP Repository	516
Downloading Remote Management AUP Files	517
Uploading Remote Management AUP Files	517
Deleting Remote Management AUP Files	517
Emergency Restore	518
User-Defined Repositories	518
Creating a User-Defined Repository	519
Retrieving Container Files	520
Uploading Files to a Repository	520
Deleting a Repository	521
Updating Repository Settings	521
Managing Files in a Repository	522
Retrieving a File from the Repository	522
Uploading a File from the Repository	522
Pre-Defined Repositories	523
Settings for Backup Files	523
Settings for Map Files	524
Settings for Parser Overrides	524
Settings for FlexConnector Files	525
Settings for Connector Properties	525
Settings for JDBC Drivers	526
Cloning Container Configuration	526
Adding Parser Overrides	527
Appendix A: Search Operators	529
cef (Deprecated)	529
chart	530
Aggregation Functions	531
Multi-Series Charts	532
The span Function	532
dedup	536
eval	536
extract	537
fields	539
head	539
keys	540
parse	541

rare	542
regex	542
rename	543
replace	544
rex	545
sort	547
tail	548
top	548
transaction	549
where	551
Appendix B: Using the Rex Operator	553
Syntax of the rex Operator	553
Understanding the rex Operator Syntax	553
Ways to Create a rex Expression	554
Creating a rex Expression Manually	555
Example rex Expressions	555
Appendix C: Logger Audit Events	559
Types of Audit Events	559
Information in an Audit Event	560
Platform Events	560
Logger Application Events	566
Appendix D: Examples of System Health Events	581
Appendix E: Event Field Name Mappings	587
Appendix F: Logger Content	593
Reports	593
Device Monitoring	594
Anti-Virus	594
CrossDevice	595
Database	599
Firewall	599
IDS-IPS	600
Identity Management	601
Network	601
Operating System	602
VPN	603
Foundation	604
Configuration Monitoring	604
Intrusion Monitoring	607
Attackers	610

Resource Access	612
Targets	614
User Tracking	616
NetFlow Monitoring	616
Network Monitoring	617
SANS Top 5	619
1 - Attempts to Gain Access through Existing Accounts	619
2 - Failed File or Resource Access Attempts	620
3 - Unauthorized Changes to Users Groups and Services	621
4 - Systems Most Vulnerable to Attack	622
5 - Suspicious or Unauthorized Network Traffic Patterns	623
Parameters	626
IPAddress	627
categoryObjectParameter	627
commonlyBlockedPorts	627
destinationAddress	628
destinationPort	628
deviceGroupParameter	628
deviceProduct	628
deviceSeverityParameter	628
deviceVendor	629
dmBandwidthParameter	629
dmConfigurationParameter	629
dmLoginParameter	629
eventNameParameter	630
resourceTypeParameter	630
webPorts	630
zoneParameter	630
zones	631
System Filters	631
Appendix G: Destination Runtime Parameters	637
Appendix H: Restoring Factory Settings	645
Before Restoring Your System	645
Restoring Your System	645
Restoring the LX500	645
Restoring LX400 and Earlier Appliance Models	647
Appendix I: Logger Search From an ArcSight Console	651
Understanding the Integrated Search Functionality	651
Prerequisites	652
Setup and Configuration	653
ESM	653

Logger	654
Supported Search Options	654
Guidelines	654
Searching on Logger From ArcSight Console	655
Index	657

Chapter 1

Overview

This chapter provides an overview of ArcSight Logger 5.5 (Logger), with references to other parts of this document for more detail.

The following topics provide an overview of Logger, including information on storage, receiver, and forwarder configuration; working with events; user management; and setup and maintenance considerations.

["Introduction" on page 21](#)
["Logger Features" on page 23](#)
["Deployment Scenarios" on page 27](#)
["Centralized Management" on page 28](#)
["What's New in Logger 5.5" on page 28](#)

Introduction

Logger is a log management solution that is optimized for extremely high event throughput, efficient long-term storage, and rapid data analysis. Logger receives and stores events; supports search, retrieval, and reporting; and can optionally forward selected events. Logger compresses raw data, but can always retrieve unmodified data on demand, for forensics-quality litigation data.

Events consist of a receipt time, event time, a source (host name or IP address), and an un-parsed message portion. Logger displays events in a tabular form, as shown in [Figure 1-1](#), adding fields that describe how Logger received the event.

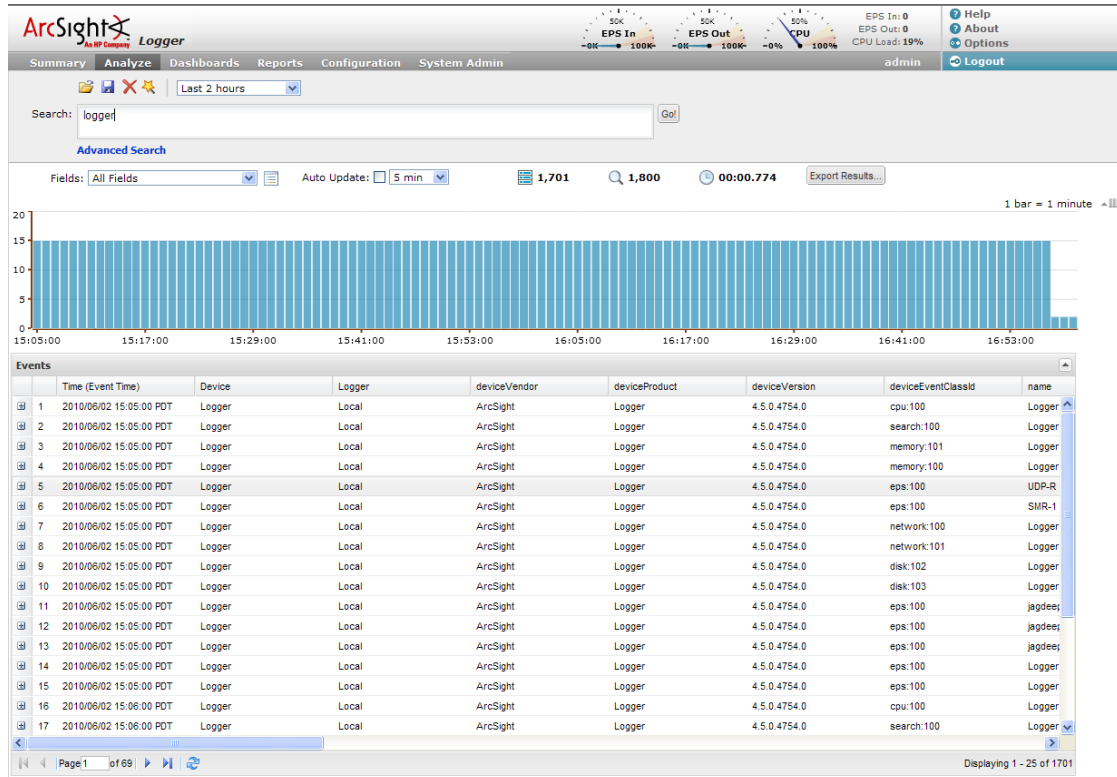


Figure 1-1 Logger web interface, Analyze tab

Similar to ArcSight Manager, Logger receives structured data in the form of normalized Common Event Format (CEF) events and unstructured data, such as syslog events. The file-type receivers configured on Logger only parse event time from an event. Although Logger is message-agnostic, it can do more with messages that adhere to the Common Event Format (CEF), an industry standard for the interoperability of event- or log-generating devices.

For more information about the Common Event Format (CEF), refer to Implementing ArcSight CEF. For a downloadable a copy of this guide, search for "ArcSight Common Event Format (CEF) Guide" on the Protect 724 Community at <https://protect724.arcsight.com>.

Logger is available in two form factors: an appliance and software. The appliance-based solution is a hardened, dedicated, enterprise-class system that is optimized for extremely high event throughput, efficient long-term storage, and rapid data analysis. The software-based solution is similar in feature and functionality to the appliance-based solution, however, the software solution enables you to install ArcSight Logger on a supported platform of your choice.

Multiple Loggers can work together to scale up to support extremely high event volume with search queries distributed across all Loggers.

Logger Features

The following sections provide an overview of key Logger features, with links to relevant sections of this guide.

Storage Configuration

The Logger appliance includes onboard storage for events. Some Logger models include RAID 1 or RAID 5 storage systems. (See Logger specifications at <http://www8.hp.com/us/en/software-solutions/enterprise-security.html>.)

On Logger appliance models that support a Storage Area Network (SAN), you need to use the SAN for storage. Logger appliance can interact with Network Attached Storage (NAS) or with a Storage Area Network (SAN) using a SAN gateway. Using a Network File System (NFS) as primary storage for events on a Logger appliance is not recommended.

On software Loggers, you need to have at least the minimum disk space described in the Release Notes to store events. The disk space needs to be on the partition where the `<install_dir>` directory exists. Specifically, most of this space should be available for the `<install_dir>/data/logger` directory. SAN can be used for storing events on both types of Loggers; however, only one LUN can be used for storing events.

On the software version of Logger, this LUN must be mapped to the `<install_dir>/data/directory` on the system on which the Logger software is installed. Using NFS as primary storage for events on the software version of Logger is not recommended.

Events are stored compressed. You cannot configure the compression level.

An NFS or a CIFS system can be used for archiving Logger data such as event archives, Saved Searches, exported filters and alerts, and configuration backup information on all Loggers. You can also configure the Logger to read event data or log files from a CIFS host.

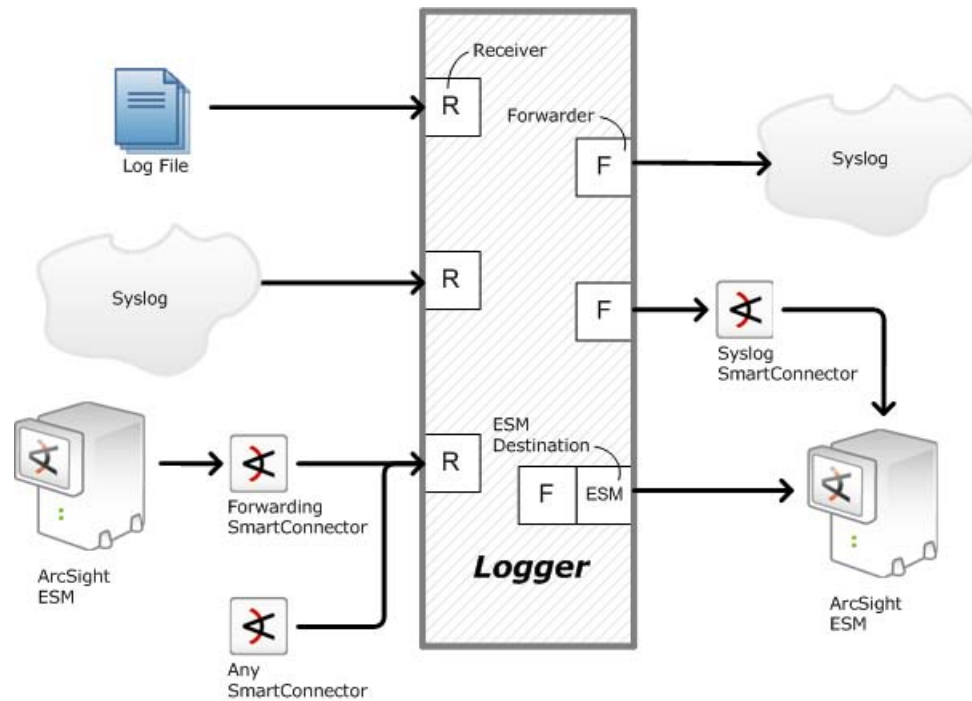
The storage volume, either external or local, can be divided into multiple storage groups, each with a separate retention policy. Two storage groups are created when Logger is first configured. New storage groups can be added later, and a storage group's size can be increased or decreased, and the retention policy defined for it can be changed.

- [“Deployment Planning” on page 29](#)
- [“Storage” on page 256](#)

Receiver Configuration

Logger receives events as syslog messages, encrypted SmartMessages, Common Event Format (CEF) messages, or by reading log files. Traditionally, syslog messages are sent using User Datagram Protocol (UDP), but Logger can receive syslog and CEF messages using the more reliable Transmission Control Protocol (TCP) as well.

Logger can also read events from text log files on remote hosts. Log files can contain one event per line or event messages that span multiple lines separated by characters such as newline (`\n`) or a carriage return (`\r`). Each event must include a timestamp. Logger can be configured to poll remote folders for new files matching a filename pattern. Once the events in the new file have been read, Logger can delete the file, rename it, or simply remember that it has been read. Logger can read remote files on network drives using SCP, SFTP, or FTP protocol, or using a previously-established NFS or CIFS mount or, on some Logger appliance models, a SAN.



Logger may also receive events from an ArcSight Manager as CEF-formatted syslog messages. These events are forwarded to Logger through a special software component called an ArcSight Forwarding SmartConnector that converts the events into CEF-formatted syslog messages before sending them to Logger.

- [“Receivers” on page 260](#)
- [“Using SmartConnectors to Collect Events” on page 52](#)
- [“Sending Events from ArcSight ESM to Logger” on page 55](#)

Analyzing Events

Events can be searched, yielding a table of events that match a particular query. Queries can be entered manually, or automatically created by clicking on terms in the event table. Queries can be based on plain English keywords (full-text search), predefined fields, or specified as regular expressions. Logger supports a flow-based search language that allows you to specify multiple search commands in a pipeline format.

By default, a Logger queries only its primary data store even if peer Loggers are configured. However, you can configure it to distribute a query across peer Loggers of your choice.

Queries can be saved as a filter or as a saved search. Saved filters can be used to select events for forwarding or to query events again later. A Saved Search is used to export selected events or save results to a file, typically as a scheduled task.

- [“Searching for Events on Logger” on page 106](#)
- [“Saving Queries \(Saved Filters and Searches\)” on page 126](#)
- [“Filters” on page 312](#)
- [“Saved Searches” on page 315](#)
- [“View and Add Parsers for Specific Log Types” on page 326](#)

Grouping Events

The combination of a source IP address and a Logger receiver is called a device. As events are received, devices are automatically created for each IP/receiver pair. Devices can also be manually created, anticipating future traffic.

Devices can be categorized by membership in one or more device groups. While an incoming event belongs to one and only one device, it can be associated with more than one device group.

Storage rules associate a device group with a storage group. Storage rules are ordered by priority, and the first matching rule determines to which storage group an incoming event will be sent.

Device groups, devices, storage groups, and peer Loggers can each be used to filter events using Search Constraints, which can be specified interactively on the Analyze page as well as when creating filters or Saved Searches.

- [“Devices” on page 247](#)
- [“Storage Rules” on page 259](#)
- [“Searching Peer Loggers \(Distributed Search\)” on page 108](#)

Exporting Events

Logger appliance can export events that match the current query locally, to an NFS mount, a CIFS mount, a SAN (on select Logger appliance models), or to the browser as a file to be downloaded.

Events from a software Logger can be exported locally to the Logger (to the <install_dir>/data/logger directory) or to the browser from which you connect to the Logger. The <install_dir>/data/logger directory can be mounted to an NFS or CIFS.

Events can be exported in Comma-Separated Values (CSV) format for easy processing by external applications or as a PDF file for generating a quick report. A PDF report includes a table of search results and any charts generated for the results. Both, raw (unstructured data) and CEF events (structured data), can be included in the PDF exported report.

Events in Common Event Format (CEF) have more columns defined, making the data more useful, but non-CEF events can be exported as well, if desired. The user can control which fields are exported.

Exports can be scheduled to run regularly by creating a Saved Search Job. First, a Saved Search is created, either manually or by saving a query on the Analyze page. A Saved Search can be based on an existing filter. A Saved Search Job combines one or more Saved Searches and a schedule with export options.

- [“Exporting Search Results” on page 119](#)
- [“Impact of Daylight Savings Time Change on Logger Operations” on page 363](#)
- [“Scheduled Saved Search” on page 316](#)

Forwarder Configuration

Logger can send events (as they are received or past events) to other hosts using UDP or TCP, to a Logger Streaming SmartConnector, or to an ArcSight Manager. The events sent to a particular host can be filtered by a query that events must match. Outgoing syslog messages can be configured to either pass the original source IP and timestamp through, or use Logger's "send time" and IP address.

Syslog messages can be sent to an ArcSight Manager using a syslog SmartConnector, but Logger can also send CEF events directly to a ArcSight Manager using a built-in SmartConnector. Logger can act as a funnel, receiving events at very high volumes and sending fewer, filtered events on to an ArcSight Manager, as shown in [Figure 1-2](#).

- ["Forwarders" on page 285](#)
- ["ESM Destinations" on page 291](#)

User Management

User accounts can be created by the Logger administrator to distinguish between different users of the system. User accounts inherit privileges from the User Group to which they belong. User Groups can have an enforced event filter applied to them, limiting the events that a specific user can see.

- ["Users/Groups" on page 392](#)
- ["Change Password" on page 407](#)
- ["Search Group Filters Tab" on page 314](#)

Other Setup and Maintenance

Logger configuration settings, such as receivers, filters, Saved Search jobs, and so on—everything except events—can be backed up as a configuration backup file to any disk and later restored.

Logs detailing Logger activity can be downloaded through the browser on demand, for debugging or other reasons. Other system information is available for viewing. Various other system settings can be modified. Some require a system reboot or restart for the changes to take effect.

Logger appliance can be rebooted using controls in the browser user interface. For the software version of Logger, the Logger service and related processes can be restarted. Follow instructions in ["Starting and Stopping the Software Logger" on page 48](#) to start, stop, or restart software Logger.

- ["Configuration Backup and Restore" on page 330](#)
- ["Retrieve Logs" on page 352](#)
- ["Storage" on page 372](#)
- ["System Locale" on page 358](#)
- ["License & Update" on page 364](#)
- ["Network" on page 358](#)

Deployment Scenarios

Typically, Logger is deployed inside the perimeter firewall with a high degree of physical security to prevent tampering with the collected event information. Logger does not require other ArcSight products. It receives and forwards syslog and log file events created by a wide variety of hardware and software network products.

Logger also interoperates with ArcSight Manager as shown in the following figures. A typical use of Logger is to collect firewall or other data and forward a subset of the data to ArcSight Manager for real-time monitoring and correlation, as shown in [Figure 1-2](#). Logger can store the raw firewall data for compliance or service-level agreement purposes.




Figure 1-2 Logger can act as a funnel, forwarding selected events to ArcSight Manager

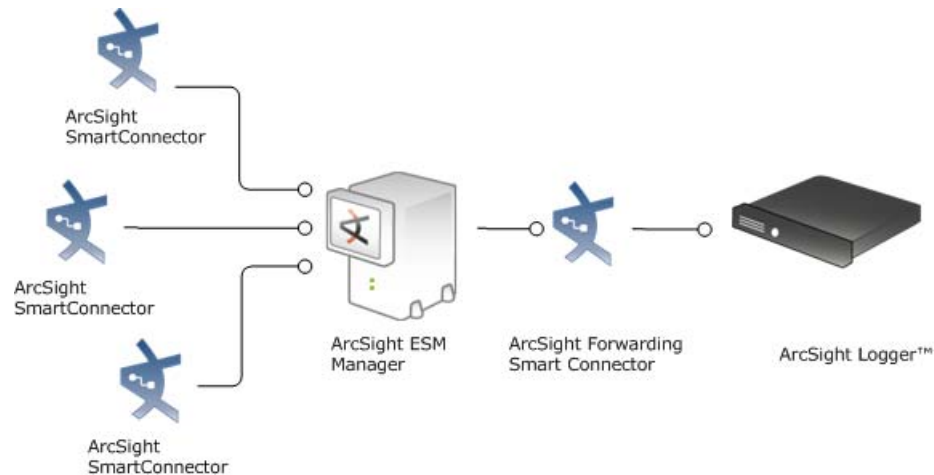


Figure 1-3 Logger can store events sent by ArcSight Manager




Figure 1-4 Logger can store and forward filtered events in a hierarchical ArcSight Manager deployment

Centralized Management

HP ArcSight Management Center (ArcMC) provides centralized management for Connector Appliances, Loggers, and software connectors with a single panel view of all managed ArcSight products.



Note

Centralized Management is not available for trial Loggers. To take advantage of this feature, you need the Enterprise version.

Using ArcSight Management Center, you can create or import configurations for managed products, and then rapidly push them to products of the same type across your network, ensuring consistent configuration for managed products with one action, and perform a variety of remote management tasks, singly and in bulk, on Connector Appliances, Loggers, and software connectors.

For more information, consult your sales representative or refer to ArcSight Management Center Administrator's Guide.

What's New in Logger 5.5

Please see the release notes for this release that are available on the HP Customer Support site (SSO) at <http://support.openview.hp.com>.

Chapter 2

Installation and Initialization

This chapter includes deployment and configuration information that is applicable to all Logger types. The installation process is specific to Logger type. Therefore, installation instructions are provided in different sections.

This chapter describes how to install and initialize the Logger appliance and software Logger. For information about installing Logger on VMWare, refer to the Quick Start Guide, for ArcSight Logger 5.5 for VMWare VM, which is available from the same location from where you download the Logger installation file.

This chapter includes information on the following topics.

- [“Deployment Planning” on page 29](#)
- [“Initial Configuration” on page 31](#)
- [“Initializing a Logger Appliance” on page 33](#)
- [“Installing a Software Logger” on page 39](#)
- [“Connecting to Logger for the first time” on page 49](#)
- [“Configuring Logger” on page 50](#)
- [“Using SmartConnectors to Collect Events” on page 52](#)
- [“Sending Events from ArcSight ESM to Logger” on page 55](#)

Deployment Planning

This section discusses the things you need to plan for before installing and initializing all Logger types. It also describes the Logger configuration the installation and initialization process sets up for you.

Storage Strategy

Logger events can be stored in these ways:

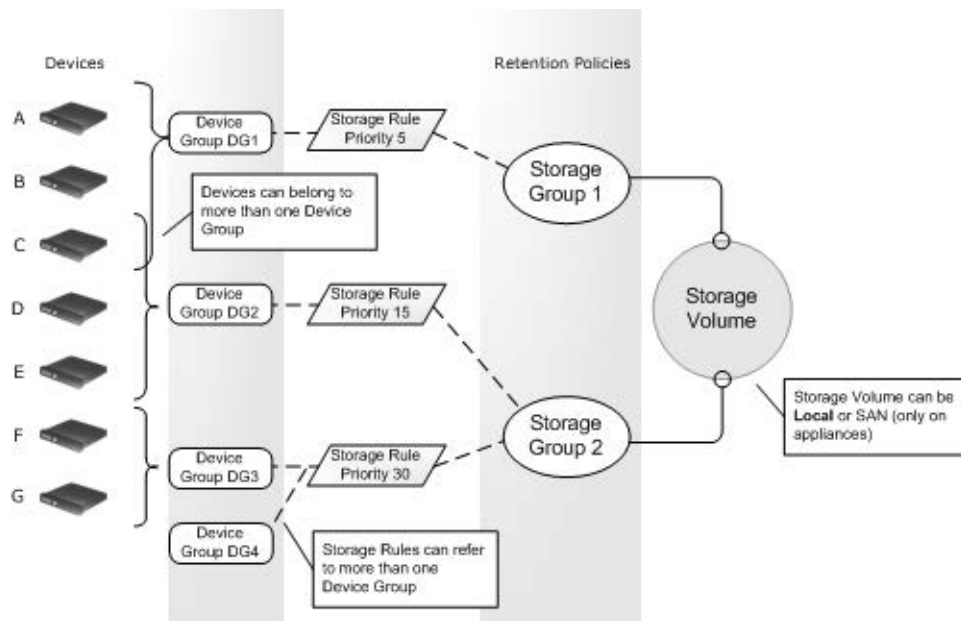
- Locally
- Remotely, on Logger appliance models that support Storage Area Network (SAN). The SAN should be available before you bring the Logger online. Only one LUN can be used for storing events.

Using a Network File System (NFS) as primary storage for events is not recommended for software Loggers. However, you can use an NFS as secondary storage for archiving data.

Retention Policy

Logger supports several storage groups, each of which can have a different retention policy. Retention policy is specified in terms of number of days that events are stored, or overall maximum size (in GB). Events from specific IP addresses can be routed to particular storage groups, making it possible to store all router events, for example, to a storage group with short retention, and business-critical host events to another storage group with a longer retention. The Logger receipt time of an event is used to determine the starting time for its retention period.

Before installing and initializing Logger, you should have an idea of your various retention policy needs, both initially and over the life span of the Logger installation.



The previous figure illustrates the relationship between ArcSight components and retention policies. Devices, on the left, are grouped by device groups. Storage groups implement different retention policies on the storage volume. Storage rules, in the middle, create a mapping between device groups and storage groups. In the example shown, Device C is a member of both Device Group 1 and Device Group 2. Storage rules are defined that send Device Group 1 events to Storage Group 1 and Device Group 2 events to Storage Group 2. There is no ambiguity, however, because each storage rule has a unique priority value, and the lower value has the higher priority. In the example, events from Device C are stored in Storage Group 1 because that storage rule has a priority of 5, which is lower than the other matching storage rule, which has a priority of 15.



Note

An implicit storage rule, with lowest priority, maps all devices to the Default Storage Group.

Initial Configuration

The installation and initialization process sets up your Logger with the initial configuration described in the sections below:

- [“SAN” on page 31](#)
- [“Storage Groups” on page 32](#)
- [“Indexed Fields and Full-text Indexing” on page 32](#)
- [“Receivers” on page 32](#)

After the initial configuration, you can do additional configuration on Logger to implement your retention policies. See [“Configuring Logger” on page 50](#) for information on devices, storage groups, and storage rules.

The initialization of a Logger appliance can only be changed by performing a factory reset (see [“Restoring Factory Settings” on page 645](#).)

SAN

This section only applies to SAN-enabled appliance models.



If you are using a SAN as your primary storage for a Logger appliance, the SAN must be set up before initializing the Logger.

By default, the HBA card on your SAN Logger has two ports. You can connect both of those ports to the same LUN for multipathing or use one port for primary storage and the other for an additional LUN for event archival, configuration backup, and export. Logger can attach to only one LUN at a time for primary storage.

When you multipath a LUN, you create two different network paths to it from Logger. Doing so reduces the possibility of a single point of failure causing the LUN to become unavailable. See [“SAN” on page 375](#) for detailed information about connecting LUN and multipathing.

Storage Volume

Logger's storage volume varies by version, up to the maximum of 4300 GB. The initialization process sets the storage volume for Trial Loggers to 7 GB. For Logger appliances, the storage volume is set to the maximum capacity for the model. For software Loggers, the storage volume is set to the maximum capacity specified in the license or the available disk space, whichever is smaller.



If Logger's maximum capacity is exceeded, events will begin to fall out of storage. For information on how to retain these events, see [“Event Archives” on page 250](#).

Storage volume can be extended after installation, but not reduced. For more information on increasing the storage volume, see [“Storage Volume Size Increase” on page 342](#).

After installing Logger, you can view the current limits on the **Configuration** (or **Configuration > Settings**) > **License Information** page. For instructions, see [“License Information” on page 351](#). For more information about licenses, including how to upload a

new one, see [“License & Update” on page 364](#) (appliance Loggers) [“License & Update” on page 414](#) (software Loggers).

Storage Groups

Two storage groups, the Default Storage Group and the Internal Event Storage Group, are created automatically during Logger initialization. These storage groups come pre-configured with the following settings:

Table 2-1 Pre-configured Default Storage Group Settings

Attribute	Appliance Logger	Software Logger
Size	Storage Volume/2	Storage Volume/2
Retention Period	180 days	180 days

Table 2-2 Pre-configured Internal Storage Group Settings

Attribute	Appliance Logger	Software Logger
Size	5 GB	3 GB
Retention Period	365 days	365 days

Logger can have a maximum of six storage groups; therefore, you can create an additional four storage groups after your Logger has been initialized. Each storage group can have different settings. You can change the retention policy and size for all storage groups, but you can only change the name of the user-defined storage groups. See [“Storage Groups” on page 257](#) for the details of adding and resizing storage groups, and changing their retention policies.

Indexed Fields and Full-text Indexing

Frequently used fields are indexed during initialization. You can add additional fields to the index, but once a field has been added, you cannot unindex it. See [“Indexing” on page 121](#) for more information. Logger comes prepared for full-text indexing.

Receivers

The default installation includes several receivers. To start receiving events, you can direct your event sources to the default receivers. After initialization, you can create additional receivers to listen for events. You can also change and delete receivers or disable and enable them as needed.

The following receivers are set up and enabled with the default installation:

- A UDP receiver—Enabled by default. The UDP receiver is on port 514 for Logger appliances. If you are installing software Logger as root, the UDP receiver is on port 514. For non-root installs, it is on port 8514. If this port is already occupied, the initialization process selects the next higher unoccupied port. This port should be allowed through any firewall rules you have configured.
- A TCP receiver—Enabled by default. The TCP receiver is on port 515 for Logger appliances. If you are installing software Logger as root, the TCP receiver is on port 515. For non-root installs, it is on port 8515. If this port is already occupied, the initialization process selects the next higher unoccupied port. This port should be allowed through any firewall rules you have configured.

- A SmartMessage receiver—Enabled by default. To receive events from a SmartConnector, download the SmartConnector and set the **Receiver Name** to be “SmartMessage Receiver” when configuring the destination.

Logger also comes pre-configured with folder follower receivers for Logger’s Apache URL Access Error log, the system Messages log, and the system Audit log (when auditing is enabled on your Linux OS). You must enable these receivers in order to use them.



Logger’s Apache URL Access Error Log, `http_error_log`, is similar in format to the Apache `access_log`. Only failed access attempts are included in the Apache URL Access Error Log.

For software Logger, the preconfigured folder follower receivers include:

- Var Log Messages—`/var/log/messages`
- Audit Log—`/var/log/audit/audit.log`
- Apache URL Access Error Log—`<install_dir>/userdata/logs/apache/http_error_log`



The folder follower receiver for the `/var/log/audit/audit.log` is only created if the folder `/var/log/audit/` already exists on your system at installation time.

Auditing is disabled on some Logger appliance models. Logger appliances that have auditing enabled will have the same preconfigured receivers as software Logger.

When auditing is disabled on the system where Logger is installed, the preconfigured folder follower receivers include:

- Var Log Messages—`/var/log/messages`
- Apache URL Access Error Log—`/opt/arcsight/userdata/logs/apache/http_error_log`

For information about receivers in general, see [“Receivers” on page 260](#).

Initializing a Logger Appliance

The information in this section describes how to configure the initial settings for your Logger appliance. It assumes that you have already installed your appliance and configured an IP address for it, as described in *Getting Started with the ArcSight Logger Appliance*, included in the shipment with your Logger appliance.

Follow these basic steps to start using your Logger appliance:

- 1 [“Acquire a License for the Logger Appliance” on page 34.](#)
- 2 [“Log In and Accept the License Agreement” on page 34.](#)
- 3 [“Initialize the Logger Appliance” on page 35.](#)
- 4 [“Set Up the Logger Appliance for Remote Access” on page 39.](#)

For information on how to install and start using your software Logger, see [“Installing a Software Logger” on page 39](#).

Acquire a License for the Logger Appliance

Starting with Logger 5.5, licenses for all Logger types are based on Daily Data (the amount of data that comes into Logger per day). The Daily Data value is monitored and enforced on software Loggers; however, currently, this value is not enforced on Logger appliances. Logger uses the sum of the sizes of the events received each day to determine this value.



Note

For software Loggers, you can increase your Daily Data limit by purchasing a higher ingestion rate in increments of 5 GB/day. While you can purchase a higher ingestion rate for software Loggers, Logger appliances come preset with the maximum ingestion capacity of the model. Therefore the ingestion capacity of Logger appliances cannot be upgraded.

A valid license file is required on the Logger appliance before you can access its functionality. If you have not obtained a license yet, follow instructions in “Hewlett-Packard Entitlement Certificate” document included in the shipment with your Logger appliance to redeem your license key. If you do not have that document, contact customer support at <https://support.openview.hp.com>.



Note

If you have multiple Logger appliances, you will need a separate license file for each of them.

After installing Logger, you can view the specific details of the current license on the License Information page (**Configuration > Settings > License Information**) and the **System Administration > License and Update page**. For more information, see “License Information” on page 351, and “License & Update” on page 364.

Log In and Accept the License Agreement

The first time you connect to the appliance through a browser, the End User License Agreement (EULA) is displayed. Before you can log in and initialize the appliance, you must review and accept the EULA.

To accept the license agreement and log in:

- 1 Use the following URL to connect to Logger through a supported browser: <https://<IP address>>, where <IP address> is the new IP address you just configured. Once you connect, the End User License Agreement is displayed.
- 2 Review and accept the license agreement. After you accept the agreement, the Login screen is displayed.

- 3 Enter your user name and password, and click **Login**.

Use the following default credentials if you are connecting for the first time or have not yet changed the default credentials:

Username: admin

Password: password



For security reasons, be sure to change the default credentials as soon as possible after connecting to Logger for the first time. Refer to [“Change Password” on page 407](#) for instructions.

- 4 Once you have successfully logged in, proceed to the section, [“Initialize the Logger Appliance” on page 35](#).

For more information about the log in screen and connecting to Logger, see [“Connecting to the Logger User Interface” on page 57](#).

Initialize the Logger Appliance

After you accept the End User License Agreement and log in for the first time, the Logger Configuration screen is displayed. On this screen, you must upload the license file and configure the initial settings for your Logger appliance. Once you complete that configuration, your Logger appliance will be ready for use.



Logger Configuration

To initialize the Logger appliance:

- 1 On the Logger Configuration screen, under **Select License File to Upload**, navigate to or specify the path and filename of the license for the Logger appliance, and click **Upload License**. If you do not have a license, see [“Acquire a License for the Logger Appliance” on page 34](#).

After the upload, the License pane displays updated license status information.

- 2 Under **System Locale Setting**, select a **Locale** for this Logger appliance from the drop-down list.

The locale setting ensures that the user interface displays information such as date, time, numbers, and messages in the format and language appropriate for the selected country. Once configured, this setting cannot be changed.

- 3 Under **Date/Time Settings**, ensure that the “Current Time Zone” and the “Current Time” settings are correct for your environment.

Click **Change Time Zone** and **Change Date/Time**, respectively, to update the time settings. For more information, see [“Time/NTP” on page 361](#).

- 4 Save your changes.

- ◆ For non-SAN models, click **Save**.

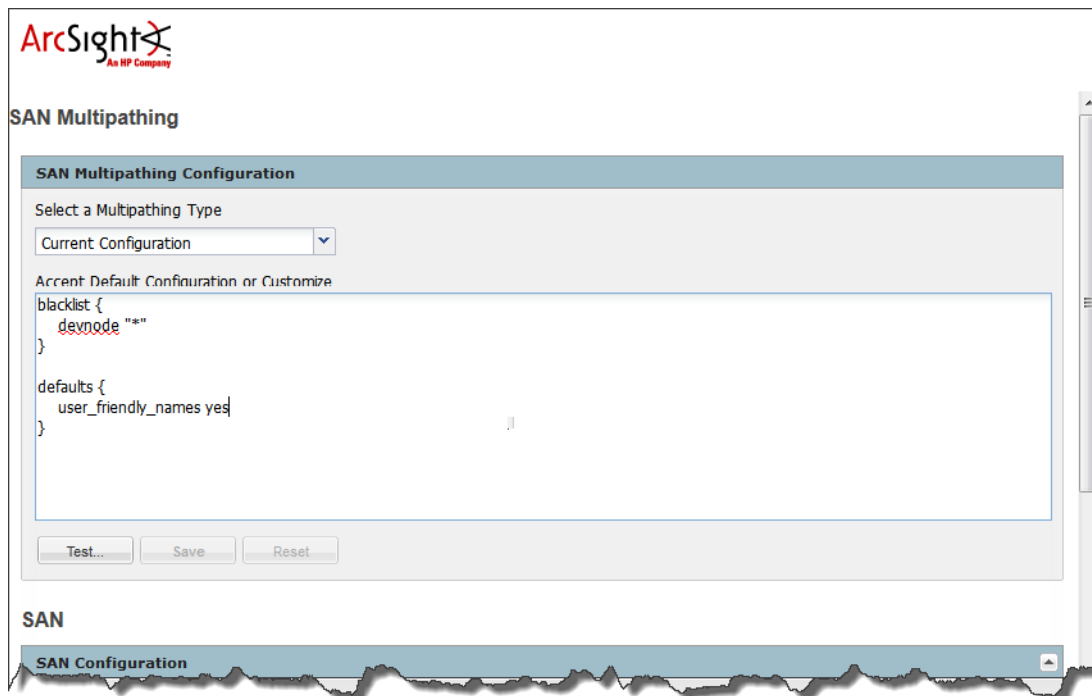
The Logger initialization process begins. Once the initialization is complete, the system reboots.

Now that you are done installing and initializing your Logger, go to the section, [“Configuring Logger” on page 50](#) for information on how to set up your Logger to start receiving events.

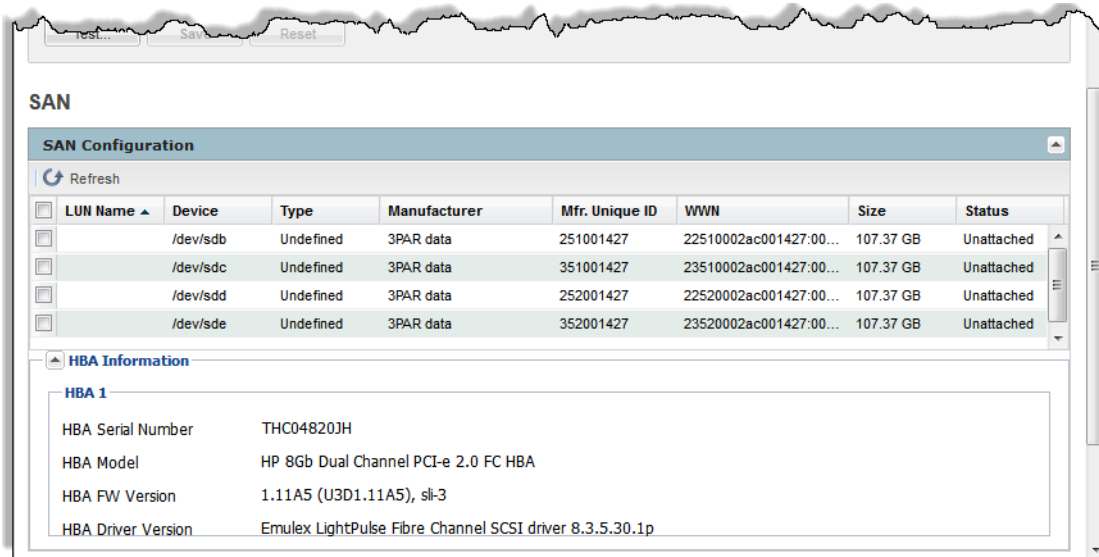
- ◆ For **SAN models**, click **Save and Configure SAN**. See [“SAN” on page 375](#) for more information about connecting LUN and multipathing.

- ◆ The SAN Multipathing Configuration screen is displayed. This screen has two panes.

The upper pane, SAN Multipathing Configuration, displays multipathing information.



The lower pane, SAN Configuration, displays the currently attached or connected LUN or LUNs.



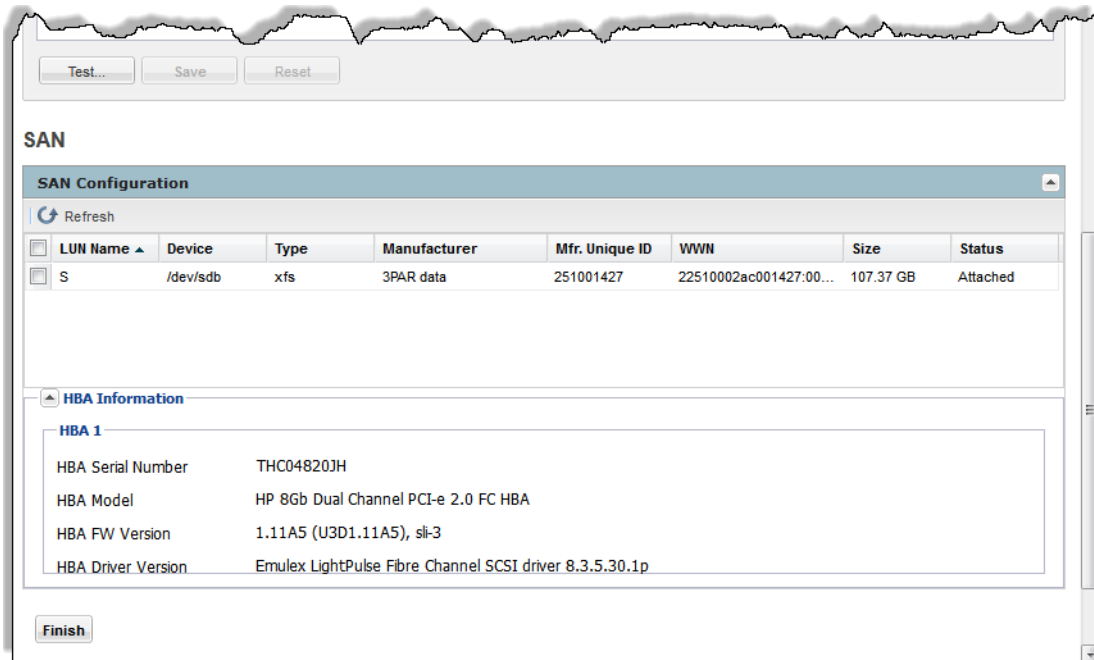
- 5 If you plan to use a single path, proceed to [Step 7 on page 38](#).
- 6 If you plan to multipath, follow these steps to configure multipathing on Logger.



If your SAN environment is set up to use multipathing, you must configure multipathing on the Logger.

- a Scroll up to the SAN Multipathing pane and select a type from the **Select a Multipathing Type** drop-down list.
- b Click **Test** to review your configuration. You cannot change the configuration once it has been saved. Be sure to review it carefully.
- c The test screen displays the routes that are currently recognized and will be multipathed if you save. When you are done reviewing the configuration, click **Close**.
- d You can accept the default configuration or customize it. Once you are satisfied with the configuration, click **Save**, and then click **OK**.
- e Click **Refresh** from the top left of the SAN Configuration pane.

The multipathed device is displayed on the SAN Configuration pane.



- 7 Select the device's checkbox in the SAN Configuration pane. The **Attach** button is displayed.



Storage volume can be extended but the size of a LUN cannot be changed after it has been attached to the Logger.

Note

For more information, see ["Storage Volume Size Increase" on page 342](#)

- 8 Click **Attach** from the top left of the SAN Configuration pane.



If no LUNs are available to attach, consult your SAN administrator to have storage allocated for the Logger.

Note

- 9 Enter a mount name for the selected LUN and click **OK**. The LUNs Attachment Status will change to "Attached" when the LUN is ready for use.

- 10 Click **Finish**.

The Logger initialization process begins. Once the initialization is complete, the system reboots.



When you configure multipath SAN connectivity to the appliance, you must also make sure that the multipathd service is configured to start on boot.

Note

Now that you are done installing and initializing your Logger, you can connect, log in, and start configuring your Logger to receive events. For instructions and information, see ["Connecting to Logger for the first time" on page 49](#) and ["Configuring Logger" on page 50](#).

Set Up the Logger Appliance for Remote Access

HP strongly recommends setting up and configuring your appliance for out-of-band remote access. Doing so ensures that you (and customer support, with your permission and assistance) can remotely access your appliance's console for troubleshooting, maintenance, and power control.

All ArcSight appliances are equipped with an HP ProLiant Integrated Lights-Out (iLO) Advanced remote management card. Follow the directions in the HP ProLiant Integrated Lights-Out User Guide to set up your appliance for remote access. The guide is available at <http://www.hp.com/go/iLO>.

Installing a Software Logger

The information in this section explains what you need to know to install and start running software Logger.

It includes information on the following topics:

- ["Downloading the Logger Software" on page 39](#)
- ["Acquiring a License for a Software Logger" on page 41](#)
- ["How Licensing Works in Software Logger" on page 40](#)
- ["Prerequisites for Installation" on page 42](#)
- ["Installation Modes" on page 42](#)
- ["Installation Steps" on page 43](#)
- ["Starting and Stopping the Software Logger" on page 48](#)
- ["Uninstalling Software Logger" on page 49](#)

For information about installing a Logger appliance, see ["Initializing a Logger Appliance" on page 33](#). For information about installing Logger on VMWare, see the Quick Start Guide for ArcSight Logger 5.5 for VMWare VM, which is available from the same location from where you download the Logger installation file.

Supported Platforms

You need to have a server with supported operating system and storage available to install the software Logger. For information about the platforms on which you can install and use Logger, refer to the Release Notes for your version.

Downloading the Logger Software

The software Logger is available in these types: Trial Version and the Enterprise Version. The Trial Version is free. Use the following table to determine where you can download the software.

Software Logger type...	Download from...
Trial Version (Free!)	HP Software Depot at http://software.hp.com
Enterprise Version	Follow the URL included in the Electronic Delivery Receipt you receive from HP in an email after placing the order.

How Licensing Works in Software Logger

Starting with Logger 5.5, licenses for all Logger types are based on **Daily Data** (the amount of data that comes into Logger per day). The Daily Data value is monitored and enforced on software Loggers; however, currently, this value is not enforced on Logger appliances. Logger uses the sum of the sizes of the events received each day to determine this value.



Note

For software Loggers, you can increase your Daily Data limit by purchasing a higher ingestion rate in increments of 5 GB/day. While you can purchase a higher ingestion rate for software Loggers, Logger appliances come preset with the maximum ingestion capacity of the model. Therefore the ingestion capacity of Logger appliances cannot be upgraded.

See [“Acquiring a License for a Software Logger” on page 41](#) for information about licensing your software Logger.



Note

If you are using ArcSight Connectors to send events to the software Logger, make sure you are running connector version 5.1.3.5870.0 or later on your connectors to ensure that event size is accurately accounted on the Logger.

Even if this limit is exceeded, the Logger continues to collect and store events; therefore, no events are lost. However, if the limit is exceeded on more than five days in a 30-day sliding window, all features involving search are disabled. **The disabled search features include forwarders as well as all searching and reporting functionality.** If this limit is exceeded six or more times (any six days or more days) in a given 30-day period, you cannot forward, search, or run reports on the collected events until the 30-day sliding window contains five or less data limit violations.

For example, you install the Logger software on January 1 with a data storage limit of 20 GB and start collecting events. Your Logger receives more than 20 GB of event data on these dates: January 5th, 13th, 18th, 19th, and 20th. Because there are five violations so far, you can forward, search, and report on the stored event data on January 21st. However, if there is another violation on January 30th, you cannot forward, search, or report on January 31st because the number of violations has exceeded the maximum allowed. (A search run on January 31st fails and the user interface displays a warning.) If there are no additional data storage-limit violations from January 31st to February 4th, the ability to forward, search, and report resumes on February 5th because the January 5th violation is now outside of the 30-day window.

The Data Volume Restrictions page (**Configuration > License Information > Data Volume Restrictions**) lists the data stored on your software Logger on day-by-day basis in the last 30 days. It also indicates the days on which data limits were exceeded, as shown in the following figure.

Summary

Analyze

Dashboards

Reports

Configuration

System Admin

admin

Logout

Devices

Event Archives

Storage

Event Input

Event Output

Alerts

Scheduled Tasks

Filters

Saved Search

Search

Peer Loggers

Configuration Backup

System Maintenance

License Information

Retrieve Logs


Content Management

License Information

Data Volume Restrictions

Date	Daily Data (MB)	Limit Exceeded
12/14/13	0	false
12/15/13	0	false
12/16/13	0	false
12/17/13	0	false
12/18/13	0	false
12/19/13	0	false
12/20/13	0	false
12/21/13	0	false
12/22/13	0	false
12/23/13	0	false
12/24/13	0	false
12/25/13	0	false
12/26/13	0	false
12/27/13	0	false
12/28/13	0	false
12/29/13	33844	true
12/30/13	0	false
12/31/13	0	false

When a data limit violation occurs, the Search user interface displays a warning, as shown in the following figure.



If you exceed the Daily Data limit frequently, you should consider purchasing a license that suits your needs. Contact your HP ArcSight sales representative to purchase a new license. Once you obtain the new license, follow the instructions in the [“License & Update” on page 414](#) to apply it on your Logger.

Acquiring a License for a Software Logger

See [“How Licensing Works in Software Logger” on page 40](#) for more information about Logger Licenses.

Software Logger includes a trial license that you can use for a limited period of time for test and evaluation purposes.

- The Enterprise Version of Logger requires a license file. You can apply the license file when you install Logger or upload one later.
- The Trial Version uses the trial license. You can upgrade by purchasing the Enterprise Version and applying the license file.

To acquire the license, follow the instructions in the Electronic Delivery Receipt you receive from HP in an email after you place the order.



Before deploying Logger on a production system, be sure to apply the license.

After installing Logger, you can view the specific details of the current license on the License Information page (**Configuration > License Information > License Information**) and the **System Administration > License and Update page**. For

more information, see [“License Information” on page 351](#), and [“License & Update” on page 414](#).

Prerequisites for Installation

Make sure these prerequisites are met before you install a software Logger:

- Before deploying in a production environment, get valid license file. If you do not have a license file, see [“Acquiring a License for a Software Logger” on page 41](#).



Note

Software Logger includes a limited trial license for test and evaluation purposes. Trial licenses do not apply to upgrades.

- You need a separate license file for each instance of software Logger. A license file is uniquely generated for each Enterprise Version download.
- Make sure a non-root user account exists on the system on which you are installing Logger.
- Decide whether to install Logger while logged in as root or as a non-root user. Your installation options vary depending on which user you choose.
 - ◆ When you install as root, a non-root user account is still required.
 - ◆ When you install as root, you can choose to configure Logger to start as a service and select the port on which Logger listens for secure web connections.
 - ◆ When you install as a non-root user, Logger can only listen for connections on port 9000. You cannot configure the port to a different value.
 - ◆ When upgrading, you cannot change the previous installation to a root-user installation. You will need to use the previously configured port 9000 for accessing software Logger.
- The hostname of the machine on which you are installing Logger cannot be “localhost”. If it is, change the hostname before proceeding with the installation.
- You must not have an instance of MySQL installed on the Linux machine on which you install Logger. If an instance of MySQL exists on that machine, uninstall it before installing Logger.
- If you want to use the GUI mode of installation and will be installing Logger software over an SSH connection, make sure that you have enabled X window forwarding using the -X option so that you can view the screens of the installation wizard. If you will be using PuTTY, you will also need an X client on the machine from which you are connecting to the Linux machine.
- Installation on 64-bit systems requires `glibc-2.12-1.25.el6.i686` and `nss-softokn-freebl-3.12.9-3.el6.i686`. Install these packages if the installation fails with the following error message, “Installation requirements not met. Pre-install check failed: 32-bit compatibility libraries not found.”

Installation Modes

The software Logger can be installed in the following three modes:

- GUI—In this mode, a wizard steps you through the installation and configuration of software Logger. For instructions, see [“Using GUI Mode to Install Software Logger” on page 43](#).

- **Console**—In this mode, a command-line process steps you through the installation and configuration of software Logger. For instructions, see [“Using the Console Mode to Install Software Logger” on page 45](#).
- **Silent**—In this mode, you provide the input required for installation and configuration through a file. Therefore, you do not need to interact with the installer to complete the installation and configuration. However, before you can use this mode, you must run the installation and configuration using one of the other modes to record the input in a file. For instructions, see [“Using the Silent Mode to Install Software Logger” on page 46](#).

Installation Steps

This section describes all three modes of software Logger installation.

Using GUI Mode to Install Software Logger

You can install software Logger as a root user or as a non-root user. See [“Prerequisites for Installation” on page 42](#) for details and restrictions.

To install the Logger software using the GUI mode:

- 1 Run these commands from the directory where you copied the Logger software:

```
chmod +x ArcSight-logger-5.5.0.XXXX.0.bin
```

```
./ArcSight-logger-5.5.0.XXXX.0.bin
```

- 2 The installation wizard launches, as shown in the following figure. Click **Next**.



You can click **Cancel** to exit the installer at any point during the installation process.



Do not use the Ctrl+C to close the installer. If you use Ctrl+C to exit the installer and then uninstall Logger, uninstallation may delete your /tmp directory.

- 3 The License Agreement screen is displayed. Scroll to the bottom of the license agreement to review the agreement and enable the “I accept the terms of the License Agreement” button.

- 4 Select **I accept the terms of the License Agreement** and click **Next**.
- 5 If Logger is currently running on this machine, an Intervention Required message is displayed. Click **Continue** to stop all current Logger processes and proceed with the installation, or click **Quit** to exit the installer.

The installer stops the running Logger processes and checks for other installation prerequisites. A message is displayed asking you to wait. Once all Logger processes are stopped and the checks complete, the next screen is displayed.

- 6 Navigate to or specify the location where you want to install Logger. By default, the /opt directory is specified.



Note

The user you are installing with must have access to the parent directory of the install directory. Otherwise, users will not be able to connect to the Logger UI and will see the following error message when they try to connect, "Error 403 Forbidden. You don't have permission to access / on this server".

- 7 If there is not enough space to install the software at the location you specify, a message is displayed. To proceed with the installation, specify a different location or make sufficient space at the location you specified. Click **Previous** to specify another location or **Quit** to exit the installer.
- 8 If Logger is already installed at the location you specify, a message is displayed. Click **Upgrade** to continue or **Previous** to specify another location. For upgrade instructions and information, refer to the Release Notes for your version.
- 9 Indicate the type of license that you want to use.
 - ◆ To evaluate Logger using the trial license, select **No, use the trial license**, and then click **Next**.
If you start with a trial license, you can upload the license file for the Enterprise Logger later. You do not need to upload a license to use the trial Logger.
 - ◆ Selecting **Yes** requires that you have already purchased the Enterprise Logger for a production environment and acquired a license file.
If you have a valid license file, select **Yes** and then click **Next**.
Click **Choose**, navigate to the license file for this software Logger, and then click **Next**.
- 10 Review the pre-install summary and click **Install**.
Installation may take a few minutes. Please wait. Once installation is complete, the next screen is displayed.
- 11 **If you are logged in as a root user** on the system on which you are installing Logger software, fill in the following fields and click **Next**.

Field	Notes
Non-root user name	This user must already exist on the system.
HTTPS port	The port number to use when accessing the Logger UI. You can keep the default HTTPS port (443) or enter any other port that suits your needs. If you specify any port except 443, users will need to enter that port number in the URL they use to access the Logger UI.

Field	Notes
Configure Logger as a service	<p>Indicate whether to configure Logger to run as a service.</p> <p>Select this option to create a service called <code>arcsight_logger</code>, and enable it to run at levels 2, 3, 4, and 5.</p> <p>If you do not enable Logger to start as service during the installation process, you still do so later. For instructions on how to enable Logger to start as a service, see “System Settings” on page 26.</p>

12 Select the locale of this installation and click **Next**.

13 Click **Next** to initialize Logger components.

Initialization may take a few minutes. Please wait. Once initialization is complete, the next screen is displayed.

14 Click **Next** to configure storage groups and storage volume and restart Logger.

Configuration may take a few minutes. Please wait. Once configuration is complete, Logger starts up and the next screen is displayed.

15 Click **Done** to exit the installer.



Note

Ensure that Logger's publicly-accessible ports are allowed through any firewall rules that you have configured.

For root installs, access to the port 443 must be allowed, plus the ports for any protocol that the logger receivers need, such as port 514 for the UDP receiver and port 515 for the TCP receiver.

For non-root installs, access to port 9000 must be allowed, plus the ports for any protocol that the logger receivers need, such as port 8514 for the UDP receiver and port 8515 for the TCP receiver.

Note: The ports listed here are the default ports. Your Logger may use different ports.

Now that you are done installing and initializing your Logger, you can connect, log in, and start configuring your Logger to receive events. For instructions and information, see [“Connecting to Logger for the first time” on page 49](#) and [“Configuring Logger” on page 50](#).

Using the Console Mode to Install Software Logger

Make sure the machine on which you will be installing the software Logger complies with the platform requirements listed in the Release Notes for your version, and that the prerequisites listed in [“Prerequisites for Installation” on page 42](#) are met.

You can install software Logger as a root user or as a non-root user. See [“Prerequisites for Installation” on page 42](#) for details and restrictions.

To install the software Logger using the Console mode:

1 Run these commands from the directory where you copied the Logger software:

```
chmod +x ArcSight-logger-5.5.0.XXXX.0.bin
./ArcSight-logger-5.5.0.XXXX.0.bin -i console
```

2 The installation wizard launches in command-line mode, as shown below. Press **Enter** to continue.

Introduction

InstallAnywhere will guide you through the installation of ArcSight Logger 5.5.

It is strongly recommended that you quit all programs before continuing with this installation.

Respond to each prompt to proceed to the next step in the installation. If you want to change something on a previous step, type 'back'.

You may cancel this installation at any time by typing 'quit'.

PRESS <ENTER> TO CONTINUE:

- 3 The next screens display license information. Installation and use of Logger 5.5 requires acceptance of the license agreement. Press Enter to display each part of the license agreement, until you reach the following prompt:

DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N):

- 4 Type Y and press Enter to accept the terms of the License Agreement.
- 5 The subsequent prompts are exactly similar to the ones described for the GUI mode install in ["Using GUI Mode to Install Software Logger" on page 43](#). Follow the instructions provided for the GUI mode install to complete the installation.

Using the Silent Mode to Install Software Logger

Before you install software Logger in silent mode, you need to create the properties file required for the silent mode installation. Once you have generated the file, you can use it for silent mode installations.

About Licenses for Silent Mode Installations

As for any Logger installation, each silent mode installation requires a unique license file. You must obtain licenses as described in ["Acquiring a License for a Software Logger" on page 41](#) and place them on the machines on which you will be installing Logger in silent mode, or ensure that the location where the licenses are placed is accessible from those machines.

Generating the Silent Install Properties File

To generate a properties file to be used for future silent installations:

- 1 Log in to the machine on which you can install software Logger to generate an installation properties file.

If you want the silent mode installations to be done as root user, log in as root. Otherwise, log in as a non-root user.

- 2 Run these commands:

```
chmod +x ArcSight-logger-5.5.0.XXXX.0.bin
```

```
./ArcSight-logger-5.5.0.XXXX.0.bin -r <directory_location>
```

where <directory_location> is the location of the directory where the generated properties file will be placed.

The properties file is called `installer.properties`. You cannot specify or change this name.

- 3 Install Logger in GUI mode, as described in [“Using GUI Mode to Install Software Logger” on page 43](#).
- 4 Once the installation completes, navigate to the directory location you specified for the `installer.properties` file earlier.

The following is an example of a generated `installer.properties` file.

```
# Fri May 11 18:27:49 PDT 2012
# Replay feature output
# -----
# This file was built by the Replay feature of InstallAnywhere.
# It contains variables that were set by Panels, Consoles or
# Custom Code.

#Choose Install Folder
#-----
USER_INSTALL_DIR=/opt/Logger/53

#License Information
#-----
LICENSE_LOCATION=/home/user/arcsight.lic
```

Installing Software Logger in Silent Mode

Make sure the machine on which you will be installing the software Logger complies with the platform requirements listed in the Release Notes for your version, and that the prerequisites listed in [“Prerequisites for Installation” on page 42](#) are met.

If you are installing as root, make sure that non-root user account that you entered when generating the silent mode properties file exists on the machines on which you are using the silent installer to install Logger.

To install the software Logger using the Silent mode:

- 1 Copy the silent mode properties file you generated previously to the same location where you have copied the Logger software.
- 2 Edit the `LICENSE_LOCATION` property in the silent mode properties file to include the location of license file for this instance of installation. (A unique license file is required for each instance of installation.)

OR

Set the `LICENSE_LOCATION` property to point to a file, such as `software_logger_license.zip`. Then, for each instance of the silent mode installation, copy the relevant license file to the location and rename it to `software_logger_license.zip`. Doing so will avoid the need to update the combined properties file for each installation.

- 3 Run these commands from the directory where you copied the Logger software:

```
chmod +x ArcSight-logger-5.5.0.XXXX.0.bin

./ArcSight-logger-5.5.XXXX.0.bin -i SILENT -f <path to
installer.properties>
```

The rest of the installation and configuration proceed silently, without requiring any input from you.

Starting and Stopping the Software Logger

The `loggerd` command enables you to start or stop the Logger software running on your machine. In addition, the command includes a number of subcommands that you can use to control other processes that run as part of the Logger software. If your Logger is installed to run as a system service, use the `service` command to start, stop, or check the status of a process on Logger.

```
<install_dir>/current/arcsight/logger/bin/loggerd
{start|stop|restart|status|quit}
```

```
<install_dir>/current/arcsight/logger/bin/loggerd {start
<process_name> | stop <process_name> | restart <process_name>}
```

```
/etc/init.d/service arcsight_logger {start | stop | status}
```

The following screen shot lists the processes that can be started, stopped, or restarted with `loggerd`.

Process Status

Refresh Status

System section

System	Status	Log	CPU Usage	Memory Usage	Data Collected
mutsum35-1007/arcsight.com	running		[0.75] [0.66] [0.58]	14.8%us 3.6%sw 1.2%wa	26.2% [1603580 KB]

NOTE: This Start/Stop buttons are for diagnostic purposes. Please use them with care.

Processes section

Process	Status	Uptime	CPU Usage	Memory Usage
apache	running	14m	14m	0.1% [7400 KB]
aps	running	14m	14m	0.2% [219336 KB]
connector	running	15m	15m	0.0% [568 KB]
insp	running	15m	15m	0.3% [19892 KB]
mysqld	running	15m	15m	0.3% [20520 KB]
postgresql	running	15m	15m	0.1% [9192 KB]
processors	running	14m	14m	0.9% [56452 KB]
receivers	running	13m	13m	0.5% [34232 KB]
reportengine	running	14m	14m	3.0% [188256 KB]

The following table describes the subcommands available with `loggerd` and their purpose.

Command	Purpose
<code>loggerd start</code>	Start all processes listed under the System and Process sections in the figure above. Use this command to launch Logger.
<code>loggerd stop</code>	Stop processes listed under the Process section only. Use this command when you want to leave <code>loggerd</code> running but all other processes stopped.

Command	Purpose
<code>loggerd restart</code>	<p>This command restarts processes listed under the Process section only.</p> <p>Note: When the <code>loggerd restart</code> command is used to restart Logger, the status message for the “aps” process displays this message:</p> <p>Process ‘aps’ Execution failed.</p> <p>After a few seconds, the message changes to:</p> <p>Process ‘aps’ running.</p>
<code>loggerd status</code>	Display the current status of all processes.
<code>loggerd quit</code>	Stops all processes listed under the System and Process sections in the figure above. Use this command to stop Logger.
<code>loggerd start <process_name></code>	Start the named process. For example, <code>loggerd start apache</code>
<code>loggerd stop <process_name></code>	Stop the named process. For example, <code>loggerd stop apache</code>
<code>loggerd restart <process_name></code>	Restart the named process. For example, <code>loggerd restart apache</code>

Uninstalling Software Logger

If you will be uninstalling the Logger software over an SSH connection in and want to use GUI mode, make sure that you have enabled X window forwarding using the `-X` option so that you can view the screens of the uninstall wizard. If you will be using PuTTY, you will also need an X client on the machine from which you are connecting to the Linux machine.

To uninstall the software Logger, enter this command in the directory where you installed the software Logger:

```
./UninstallerData/Uninstall_ArcSight_Logger_5.5
```

The uninstall wizard is launched. Click **Uninstall** to start uninstalling Logger.

Connecting to Logger for the first time

The Logger user interface is a web browser application using Secure Sockets Layer (SSL) encryption. Users must be authenticated with a name and password before they can use the interface. Refer to the Release Notes document to find out the browsers and their versions supported for this release.

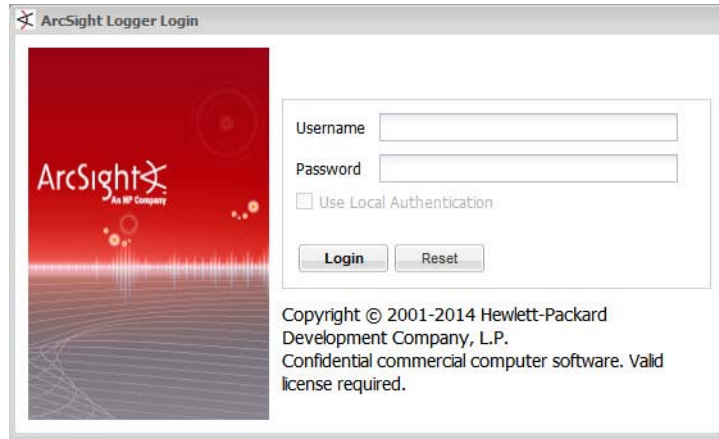
To connect and log into Logger:

- 1 Use the following URL to connect to Logger through a supported browser:
`https://<hostname or IP address>:<configured_port>`

where the `hostname or IP address` is the system on which the Logger software is installed, and `configured_port` is the port specified during the Logger installation.

Once you connect, the following Login screen is displayed.

- 2 Enter your user name and password, and click **Login**.



Use the following default credentials if you are connecting for the first time or have not yet changed the default credentials:

- 3 Username: admin
Password: password
- 4 Once you have successfully logged in, go to the section, [“Configuring Logger” on page 50](#) for information on how to set up your Logger to start receiving events.



For security reasons, be sure to change the default credentials as soon as possible after connecting to Logger for the first time. Refer to [“Change Password” on page 407](#) for instructions.

For more information about the log in screen and connecting to Logger, see [“Connecting to the Logger User Interface” on page 57](#).

Configuring Logger

Once you have logged in successfully, you can enable the preconfigured receivers and configure devices, device groups, and storage groups necessary to implement your retention policy. If you have installed multiple Loggers, you must connect to each and configure it separately.

Receivers

Now that you have finished installing Logger, you can set up receivers to listen for events. You can use the preconfigured receivers or add your own. Logger comes with preconfigured with a TCP receiver, a UDP Receiver, and a SmartMessage receiver already enabled and ready to receive events. Logger also comes pre-configured with folder follower receivers for Logger’s Apache Access Error Log, the system Messages Log, and the system Messages Audit Log (if auditing is enabled on your Linux OS). You must enable these receivers in order to use them. Receivers can be disabled and re-enabled later. You can add, change, and delete them as needed.

For more information on receivers, see [“Receivers” on page 260](#).

Enabling the Pre-configured Receivers

When you first log in by using the URL you configured, Logger will display a banner like the one below, telling you about the disabled receivers.



Click the link in the banner to open the Receivers page.



On software Loggers:

Before enabling these receivers, you must make `/var/log/audit/audit.log` and `/var/log/messages` readable by the non-root user you installed with or specified during Logger installation.

Receivers

Source Types

Parsers

Add

Once you enable the Apache URL Access Error Log receiver, Logger will start storing entries from the `<install_dir>/userdata/logs/apache/http_error_log` file.

Logger can also store entries from the messages and audit.log files in the `/var/log/*` folders. Before enabling the receivers for these files, consult the [Logger Administrators guide](#) for details.

Name	Type	IP Address	Port			
Apache URL Access Error Log	Folder Follower Receiver					
Audit Log	Folder Follower Receiver					
Var Log Messages	Folder Follower Receiver					
SmartMessage Receiver	SmartMessage Receiver					
TCP Receiver	TCP Receiver	All	515			
UDP Receiver	UDP Receiver	All	514			

To enable a receiver, click the disabled icon () at the end of the row.

Alternately, you can navigate to the Receivers page from the menu to enable the receivers.

To open the Receivers page from the menu and enable a receiver:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Receivers** tab (right panel).
- 3 Click the disabled icon () at the end of the row.

For information on how to use the preconfigured SmartMessage receiver, see ["Using SmartConnectors to Collect Events" on page 52](#). For more information on enabling and disabling receivers, see ["Working with Receivers" on page 263](#).

Devices

Logger begins storing events when an enabled receiver receives data or, in the case of a file receivers, when the files become available. Using a process called autodiscovery, Logger automatically creates resources called devices to keep track of source IP addresses and uses DNS to map them to hostnames. Eventually, a device is created for each device from which Logger received events.

You can also create devices preemptively, by entering the IP addresses or hostnames of data sources that you expect to be sending events to Logger. You might do this if you do not want to wait for autodiscovery, or if you want to control the initial naming of each device. Discovered devices are named for their host, or if the DNS lookup fails, for their IP address, and their receiver. For information about creating devices, see [“Devices” on page 247](#).

Device Groups

Device groups are containers or logical groupings for devices, in the same way folders (or directories) contain files. They are a name for a group of devices. A given device can be a member of several device groups. Each device group can be associated with particular storage group, which would assign a retention policy.

You can change and delete device groups freely as your needs change. Setting up device groups initially is not critical; incoming events that are not assigned to a device group are automatically sent to the Default Storage Group. For the details of setting up device groups, see [“Device Groups” on page 249](#).

Storage Rules

Events are stored in the Default Storage Group unless otherwise specified. Storage rules are a way to direct events from certain device groups to certain storage groups. You can use them to implement additional retention policies.

If you created additional storage groups, and want to send events to them, you can do that with storage rules. If you choose not to create storage rules, events from all devices will be sent to the Default Storage Group and use its specified retention policy.

If you want to implement multiple retention policies, you can create storage rules that associate the specific device groups with the storage groups that implement the desired retention policy.

For example, you could create one device group for each retention policy. However, for more control, you could associate device groups with storage groups and storage rules and use them to categorize events. For example, you could search for events that match a certain pattern and which belong to a particular device group, and send them to a particular storage group for retention based on event category.

See [“Storage Rules” on page 259](#) for more information.

Storage rules are evaluated in order of priority; the first matching rule determines to which storage group an event is sent. This approach means that a single device can belong to several device groups without ambiguity about which storage group it will end up in.

Using SmartConnectors to Collect Events

Similar to ArcSight Manager, Logger leverages the ArcSight SmartConnectors to collect events. SmartConnectors can read security events from heterogeneous devices on a network (such as firewalls and servers) and filter events of interest (and optionally aggregate them) and send them to a Logger receiver. Logger can receive structured data in the form of normalized Common Event Format (CEF) events from the SmartConnectors.

SmartMessage

SmartMessage is an ArcSight technology that provides an efficient secure channel for Common Event Format (CEF) events between ArcSight SmartConnectors and Logger.



Caution

SmartMessage and FIPS require SmartConnector 4.7.5 or later. If you do not have the current build, download the latest from the HP ArcSight web site.

Older SmartConnectors will work with Logger, but may not support SmartMessage or FIPS.

SmartMessage provides an end-to-end encrypted secure channel using secure sockets layer (SSL). One end is an ArcSight SmartConnector, receiving events from the many devices supported by ArcSight SmartConnectors. The other end is a SmartMessage receiver on Logger.



Note

The SmartMessage secure channel uses SSL protocol to send encrypted events to Logger. This is similar to, but different from, the encrypted binary protocol used between SmartConnectors and ArcSight Manager.

Downloading SmartConnectors

For Logger appliance and the Enterprise Version of Software Logger, contact your HP ArcSight sales representative or customer support for the location to download SmartConnectors.

A restricted set of ArcSight SmartConnectors are supported and available For the Trial Version of Software Logger. You can download these SmartConnectors from the same location from which you downloaded Logger. The configuration guides for the supported SmartConnectors are available at the same web site. To learn more about ArcSight ArcSight SmartConnectors, visit

<http://www8.hp.com/us/en/software-solutions/enterprise-security.html>.

Configuring a SmartConnector to Send Events to Logger

Logger comes pre-configured with a SmartMessage Receiver. To use it to receive events from a SmartConnector, you must configure the SmartConnector as described below. You can also create new SmartMessage receivers and configure the SmartConnectors with these newly created receivers. When configuring a SmartConnector, be sure to specify the correct receiver name.

To configure a SmartConnector to send events to Logger:

- 1 Install the SmartConnector component using the SmartConnector User's Guide as a reference. Specify Logger as the destination instead of ArcSight ESM or a CEF file.



Note

If you are using the Trial Version of Logger, refer to the documentation that came with your SmartConnector for instructions.

- 2 Specify the required parameters. Enter the Logger hostname or IP address and the name of the SmartMessage receiver. These settings must match the receiver in Logger that listen for events from this connector.

- ◆ To use the preconfigured receiver, specify “SmartMessage Receiver” as the **Receiver Name**.
- ◆ To use SmartMessage to communicate between an ArcSight SmartConnector and a Logger appliance, configure the SmartConnector to use port 443.
- ◆ To communicate between an ArcSight SmartConnector and software Logger, configure the SmartConnector to use the port configured for the software Logger.
- ◆ For un-encrypted CEF syslog, enter the Logger hostname or IP address, the desired port, and choose UDP or TCP output.

Configuring SmartConnectors to Send Events to Both Logger and an ArcSight Manager

You can configure a SmartConnector to send CEF syslog output to Logger and send events to an ArcSight Manager at the same time.

For more information about the Common Event Format (CEF), refer to Implementing ArcSight CEF. For a downloadable a copy of this guide, search for “ArcSight Common Event Format (CEF) Guide” on the Protect 724 Community at <https://protect724.arcsight.com>.

- 1 Install the SmartConnector normally. Register the SmartConnector with a running ArcSight Manager and test that the SmartConnector is up and running.
- 2 Start the SmartConnector configuration program again using the `$ARCSIGHT_HOME/current/bin/runagentsetup` script (or `arcsight agentsetup -w`).
- 3 Select **I want to add/remove/modify ArcSight Manager destinations**, then choose **Add new destination**.
- 4 Choose Logger and specify the requested parameters. Restart the SmartConnector for changes to take effect.

Configuring SmartConnectors for Failover Destinations

SmartConnectors can be configured to send events to a secondary, failover, destination when a primary connection fails.

To configure a failover destination, follow these steps:

- 1 Configure the SmartConnector for the primary Logger as described above. The transport must be raw TCP in order to detect the transmission errors that trigger failover.
- 2 Edit the `agent.properties` file in the directory `$ARCSIGHT_HOME/current/user/agent`, where `$ARCSIGHT_HOME` is the root directory where the SmartConnector component was installed. Add this property:

`transport.types=http,file,cefsyslog`

Delete the `transport.default.type` property.
- 3 Start the SmartConnector configuration program again using the `$ARCSIGHT_HOME/current/bin/runagentsetup` script (or `arcsight agentsetup -w`).
- 4 Choose **I want to add/remove/modify** and, with the primary Logger selected, choose **Modify**. Then select **Add failover destination**.
- 5 Enter information for the secondary Logger.
- 6 Restart the SmartConnector for the changes to take effect.

- 7 For more information about installing and configuring ArcSight SmartConnectors, refer to the ArcSight SmartConnector User's Guide, or specific SmartConnector Configuration Guides, available from the Protect 724 Community at <https://protect724.arcsight.com>.

Sending Events from ArcSight ESM to Logger

The ArcSight Forwarding SmartConnector can read events from an ArcSight Manager and forward them to Logger as CEF-formatted syslog messages.



Note

The Forwarding SmartConnector is a separate installable file, named similar to this:

ArcSight-4.x.x.<build>.x-SuperConnector-<platform>.exe

Use build 4810 or later for compatibility with Logger.

To configure the ArcSight Forwarding SmartConnector to send events to Logger:

- 1 Install the SmartConnector component normally, but cancel the installation when the SmartConnector Wizard asks whether the target Manager uses a demo certificate.



Figure 2-1 SmartConnector Configuration Wizard

When the first screen of the SmartConnector Configuration Wizard appears, asking about a demo certificate, click **Cancel**.

- 2 Confirm that you want to exit, then click **Done** to dismiss the Install Wizard. This will install the SmartConnector, but leave it un-configured.
- 3 Create a file called **agent.properties** in the directory `$ARCSIGHT_HOME/current/user/agent`, where `$ARCSIGHT_HOME` is the root directory where the SmartConnector component was installed. This file should contain a single line:


```
transport.default.type=cefsyslog
```
- 4 Start the SmartConnector configuration program again using the `$ARCSIGHT_HOME/current/bin/runagentsetup` script (or `arcsight agentsetup -w`).

- 5 Specify the required parameters for CEF output. Enter the desired port for UDP or TCP output. These settings will need to match the receiver you create in Logger to listen for events from ArcSight ESM.

Parameter	Description
Ip/Host	IP or host name of the Logger
Port	514 or another port that matches the receiver
Protocol	UDP or Raw TCP
ArcSight Source Manager Host Name	IP or host name of the source ArcSight Manager
ArcSight Source Manager Port	8443 (default)
ArcSight Source Manager User Name	A user account on the source Manager with sufficient privileges to read events
ArcSight Source Manager Password	Password for the specified Manager user account
SmartConnector Name	A name for the ESM to Logger connector (visible in the Manager)
SmartConnector Location	Notation of where this connector is installed
Device Location	Notation of where the source Manager is installed
Comment	Optional comments

To configure the Forwarding SmartConnector to send CEF output to Logger and send events to another ArcSight Manager at the same time, see ["Configuring SmartConnectors to Send Events to Both Logger and an ArcSight Manager"](#) on page 54.

For more information about the Common Event Format (CEF), refer to Implementing ArcSight CEF. For a downloadable a copy of this guide, search for "ArcSight Common Event Format (CEF) Guide" on the Protect 724 Community at <https://protect724.arcsight.com>.

Chapter 3

User Interface and Dashboards

This chapter provides an overview of the layout of the Logger user interface. Additionally, the chapter describes dashboards available on Logger that you can use to view summarized event information, create your own dashboards for an all-in-one view of Logger information that is of interest to you, and monitoring dashboards that display the real-time and historical status of receivers, forwarders, and storage, CPU, and disk usage statistics.

This chapter includes information on the following topics.

Logging in: see [“Connecting to the Logger User Interface” on page 57](#)

Navigation: [“Navigating the User Interface” on page 59](#)

Summary Dashboard: [“Summary” on page 61](#)

Dashboards: see [“Dashboards” on page 64](#)

Performance monitoring: see [“The Default Monitor Dashboard” on page 70](#)

Connecting to the Logger User Interface

Logger works with most browsers, including Firefox and Internet Explorer. JavaScript and cookies must be enabled. An Adobe Flash Player plug-in is required for Internet Explorer browsers that access the Logger user interface. Some redundant monitoring features will be unavailable if the Flash Player plug-in is not installed. The Flash Player plug-in is available for free at <http://www.adobe.com/products/flashplayer/>.

Refer to the Release Notes document to find out the browsers and their versions supported for this release.

Logging In

The Logger user interface is a web browser application using Secure Sockets Layer (SSL) encryption. Users must be authenticated with a name and password before they can use the interface. Refer to the Release Notes document to find out the browsers and their versions supported for this release.

To connect and log into Logger:

- 1 Use the URL configured during Logger installation to connect to Logger through a supported browser.

On the Logger appliance, connect using this URL:

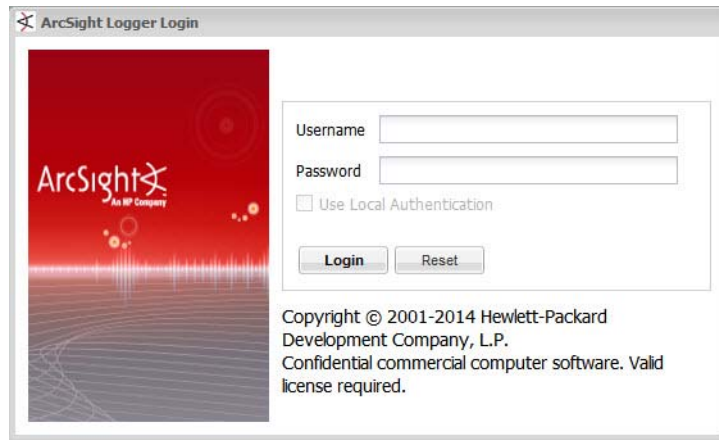
`https://<hostname or IP address of the Logger appliance>`

On software Logger, connect using this URL:

`https://<hostname or IP address>:<configured_port>`

where the hostname or IP address is the system on which the Logger software is installed, and configured_port is the port set up during the Logger installation, if applicable. (A port is not required if the installation was done as the root user.)

Once you connect, the following Login screen is displayed.



- 2 Enter your user name and password, and click Login. Use the following default credentials if you are connecting for the first time or have not yet changed the default credentials:

Username: admin

Password: password



Caution

For security reasons, be sure to change the default credentials as soon as possible after connecting to Logger for the first time. Refer to [“Change Password” on page 85](#) for instructions.

If login succeeds, the Logger user interface is displayed. If login fails, the message Authentication Failed is displayed at the top of the login screen. Enter the correct username and password combination to try again.

Depending on your system administration settings, the following options maybe also be available.

- ◆ Forgot Password?—A “Forgot Password?” link is displayed if your Logger is configured to show it. For more information on the Forgot Password link, see [“Forgot Password” on page 396](#).
- ◆ Use Local Authentication—The “Use Local Authentication” checkbox is always displayed, but only becomes active when a login attempt fails. By default, this option is available only for the default admin. For more information on the Use Local Authentication, see [“Local Password Fallback” on page 401](#).

Navigating the User Interface

As shown in [Figure 3-1](#), a navigation and information band runs across the top of every page in the user interface.

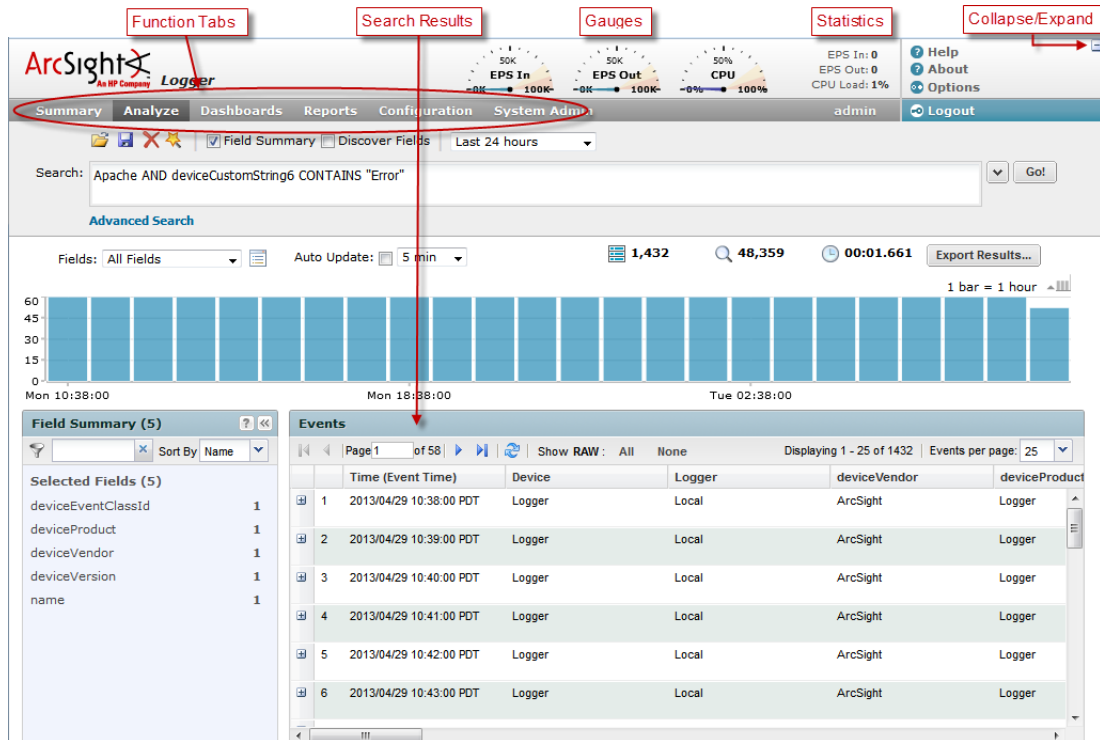




Figure 3-1 Overall layout of the Logger interface

The upper right corner of the screen includes the [Help](#), [Options](#), [About](#), and [Logout](#) links.

Gauges at the top of the screen provide an indication of the throughput and CPU usage information, available in more detail on the Monitor Dashboard ([“Dashboards” on page 64](#)). The range of the gauges can be changed on the Options page. The current logged-in user's name is shown below the statistics. The gauge and logo bar can be collapsed to allow more room on the screen for search results and reports. Click the  icon to collapse the bar, and the  icon to expand it.

The Summary, Analyze, Dashboards, Reports menu tabs provide access to various Logger functions and data stored on it. The Configuration and System Admin menus are used for configuring the system administration and configuration settings on Logger. For more information on each, refer to the sections below.

- The options available in the Summary menu are discussed in [“Summary” on page 61](#).
- The options available in the Dashboards menu are discussed in [“Dashboards” on page 64](#).
- The options available in the Analyze menu are discussed in [“Searching and Analyzing Events” on page 75](#).

- The options available in the Reports menu are discussed in [“Reporting” on page 137](#)
- The options available in the Configuration menu are discussed in [“Configuration” on page 247](#).
- The options available in the System Admin menu are discussed in [“System Admin - Logger Appliance” on page 357](#) and [“System Admin - Software Logger” on page 413](#).

Help

Click the Help link on any page to display online help for the current page. Along with other Logger documentation, the -ArcSight- Logger Administrator’s Guide is available as an Adobe Acrobat PDF document, through the -ArcSight- Product Documentation community on <https://protect724.arcsight.com>.

Options

The Options page, as shown in the following figure, allows you to set the range on the EPS In and EPS Out gauges. If the event rate exceeds the specified maximum, the range is automatically increased.

Click the Options link to set the default start page (home page) for all users and specific start pages for individual users. These pages indicate which user interface page is displayed after a user logs in.

Options

System

EPS input rate gauge max100K

EPS output rate gauge max100K

Default start page for all usersDashboards

Personal

Default start page for adminUse default for all users

SaveCancel

Use the following table to figure out how to configure a specific start page.

If you want to set...	Configure the...
The same start page for all users	Default start page for all users option to the desired value. This is a global setting for your Logger. To override this setting, configure a different start page for specific users by using the Default start page for <username> option. When you set Default start page for all users option to Dashboards, the Monitor Dashboard is the default dashboard displayed for all users, unless the users have configured other dashboards as their defaults, as described in “To set a dashboard as default:” on page 67 .

If you want to set...	Configure the...
A different start page for specific users	<p>Default start page for <username> option to the desired value.</p> <p>This setting overrides the global Default start page for all users setting.</p> <p>When this option is set to "Use default for all users", the global default page (Default start page for all users) value is used for all users.</p>
A specific dashboard for a specific user OR A specific dashboard for all users	<p>Default start page for <username> option to Dashboards.</p> <p>The Monitor Dashboard is the default dashboard displayed for all users. However, if you want to display a different dashboard for one or more users, set the desired dashboard as the default when logged in as those users. For details, see "To set a dashboard as default:" on page 67.</p>

About

Click the About link to get version information about your Logger.

Logout

Click the Logout link on any page to return to the Login screen. Logging out is good security practice, to eliminate the chance of unauthorized use of an unattended Logger session.

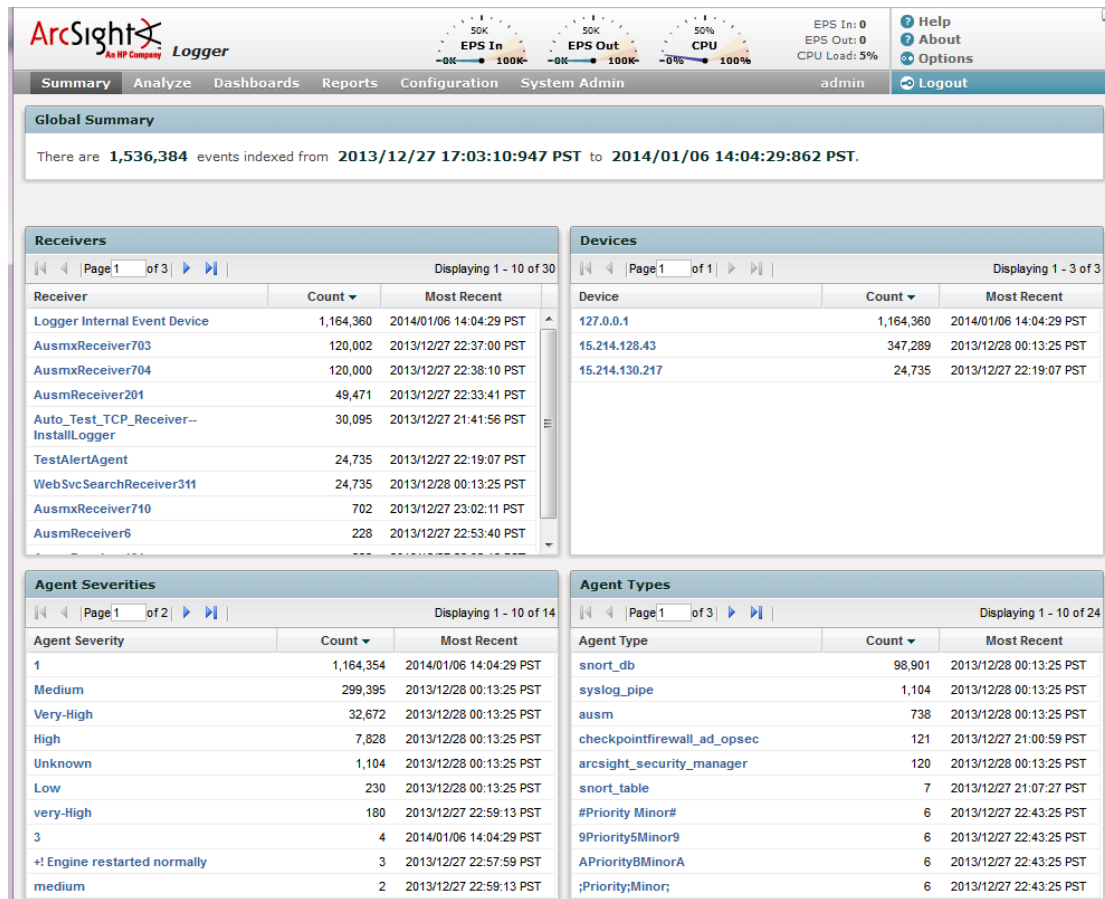
Logger automatically logs you out after a user-configurable length of time (15 minutes by default). To change this length of time, see ["Users/Groups" on page 392](#).



Simply closing the browser window does not automatically log you out. Click the Logout link to prevent the possibility of a malicious user restarting the browser and resuming your Logger session.

Summary

The Summary page is a global dashboard that provides summarized event information about your Logger in one screen. It enables you to gauge incoming events activity and the status of indexing. The events that are in Logger's primary storage (not aged out due to retention or archived data) are used to generate the summary information.



Specifically, the Summary page contains the following panels:

- Global Summary

The number of events indexed on your Logger during the time period displayed on the screen. This time period is dependent on the retention policy set on your Logger. The start is the time of the oldest event stored in the Logger since the Logger was restarted, that has not aged out due to retention; the end time is current time.

- Receivers

The list of receivers configured on your Logger, the number of events received on each receiver (that are in Logger's primary storage, not aged out due to retention or archived data), and the timestamp of the last event received on each receiver.

If a receiver is deleted, the summary information for it will continue to display until the events received on it age out from Logger's primary storage. However, the receiver name is changed to the receiver ID (a numerical string) associated with the deleted receiver.

- Devices

A device is a named event source, comprising of an IP address (or hostname) and a receiver name.

The Devices panel lists devices configured on your Logger, the number of events received on each device (that are in Logger's primary storage, not aged out due to retention or archived data), and the timestamp of the last event received on each device.

If a device is deleted, the summary information for it will continue to display until the events received on it age out from Logger's primary storage. However, you cannot click the device name to view the events associated with the deleted device.

■ Agent Severities

The list of severity levels of the incoming events from ArcSight SmartConnectors to your Logger, the number of events received of each severity level, and the timestamp of the last event received of each severity level.



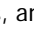

Only events in Logger's primary storage (not aged out due to retention or archived data) are considered when summarizing this information.

■ Agent Types

The list of ArcSight SmartConnectors sending events to your Logger, the number of events received from each SmartConnector (for events that are in Logger's primary storage, not aged out due to retention or archived data), and the timestamp of the last event received from each SmartConnector.

If a SmartConnector is deleted, the summary information for it will continue to display until the events received from it age out from Logger's primary storage.

The Summary page is pre-designed to display the information described above. You cannot change or add other panels to it. If you need to display other information, you can create a custom Dashboard as described in ["Dashboards" on page 64](#). The information displayed on the Summary page is for your local Logger only, and does not include information about peer Loggers even if peers are configured.

Each panel displays up to 10 items. If there are more than 10 items, click the  icon to see the additional items, and  to go to the end of the list. Similarly, click  to go to the previous 10 items, and  to go to the first 10 items.

You can drill down to view the events by a specific resource—receiver, device, agent severity, or agent type. To do so, click on the resource to go to the Analyze screen.



The Search box is automatically populated with the information you had clicked on the Summary page, and the Start and End fields are populated with the time of oldest events stored on your Logger (that have not aged out due to retention) and the current time, respectively. Click Go to run the query to search for events matching this criteria. You can further refine the search query to filter the search results to suit your needs. A drill-down on sourceType and sourceData fields is not supported.

Search Group filters that enforce privileges on storage groups are applied to the content displayed on the Summary page. However, Search Group filters that enforce privileges on *device groups* are not applied. Therefore, the Summary page includes counts of events in device groups to which a user does not have privileges. However, if the user tries to drill-down to view events, search results in accordance with access privileges are returned as the search query is run on the Analyze page, which enforces all types of Search Group filters. Similarly, if a Search Group filter enforces privileges on both, storage groups and device groups, only the storage group enforcement is applied on the Summary page.

Dashboards

Dashboards are an all-in-one view of the Logger information of interest to you. You can assemble various search queries that match events of interest to you, status of Logger resources such as receivers, forwarders, storage, CPU, and disk, or a combination of both on a single dashboard for status at-a-glance.

Each Dashboard contains one or more panels of these types:

- Search Results

The Search Results panels display events that match the query associated with the panel.

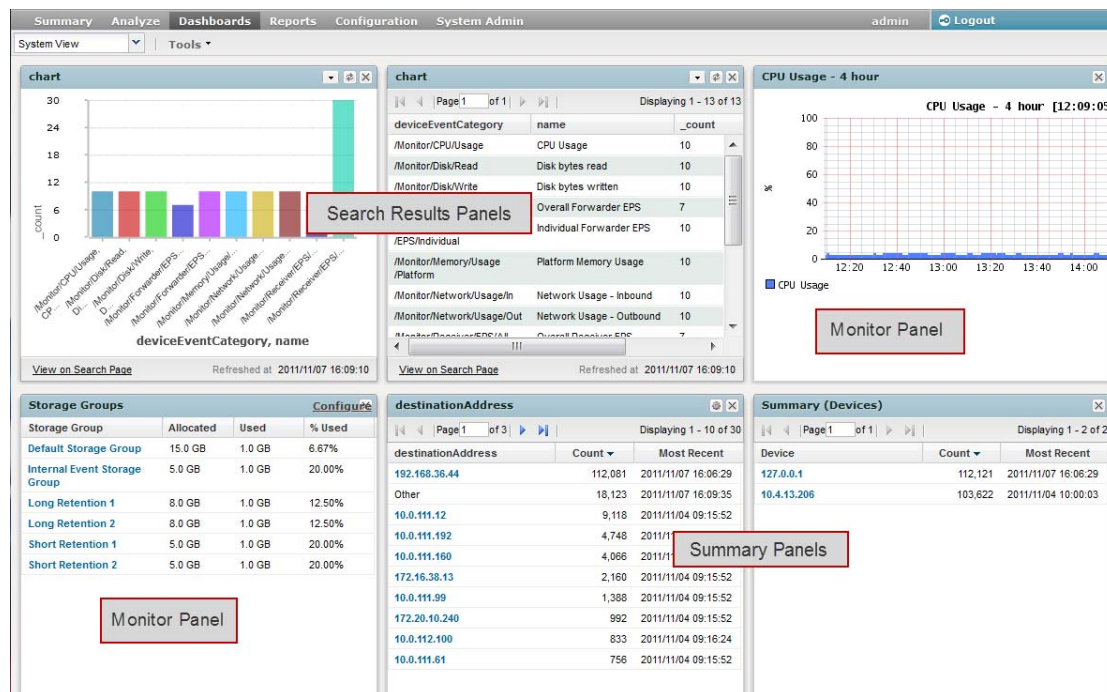
- Monitor

The Monitor panels display the real-time and historical status of various Logger components such as receivers, forwarders, storage, CPU, and disk.

- Summary

The Summary panels display summarized event information about your Logger—the number of events received of a specific resource or field type, and the timestamp of the last event received for that resource or field type.

A dashboard can contain a mix of Search Results, Monitor, and Summary panels. There is no limit on the number of Monitor and Summary panels you add to a single dashboard; however, you can only add up to four Search Results panels for optimum performance.



Each Search Results panel is associated with a saved search query. You can only associate saved search queries that contain an aggregation operator such as `chart` or `top` for this type of panel. The Search Results panel can be of two types: Search Results Chart and Search Results Table. The Search Results Chart panel displays search results in a chart form, and the Search Results Table panel displays search results in a table form, as shown in the following figure.

Type “Summary (Receivers)”. Besides the four Summary panels (Agent Severities, Agent Types, Receivers, and Devices), you can also create a user-defined Summary panel in which you can select *any indexed, non-time field* by which you want to categorize event summary. For example, if you want to add a Summary panel to display event summary categorized by “destinationAddress”, you can add a panel of Type “Summary (User Defined)” for this field if it is indexed on your Logger.

You can also drill-down on any of the resources listed in the Monitor and Summary panels you add to view events by a specific resource or field value on the Analyze (Search) page. For example, you can click on a storage group in a Monitor panel to view its events in the last 24 hours, or you can click on an event name “Network Usage - Inbound” to view all events of that name in the last one hour. Additionally, you can access the Configuration page for any of the resources listed in the Monitor panels to configure them. For example, if you want to configure a receiver, click the Configure link on top of the Monitor (Receiver) panel.

Search Group filters that restrict privileges on device groups are not enforced on *Summary panels*. Therefore, Summary panels include counts of events in device groups to which a user does not have privileges. However, if the user tries to drill-down to view events, search results in accordance with access privileges are returned as the search query is run on the Analyze page, which enforces all types of Search Group filters. Similarly, if a Search Group filter enforces privileges on both, storage groups and device groups, only the storage group enforcement is applied on Summary panels.

Users can create both shared and private dashboards.

- Shared dashboards are visible to all users with the appropriate privileges.
- Private dashboards are visible only to the creator or users with “admin” privileges.
- Only the creator or users with “admin” privileges can edit or delete dashboards of either type.

A user accessing a shared dashboard must have privileges to view the information displayed in the dashboard; otherwise, the information to which they do not have the privileges is not displayed, and the associated panel displays a message that indicates the reason for the undisplayed information.

Creating and Managing Dashboards

You need these privileges (in the Logger Rights group) to perform dashboard operations:

- “Use and view dashboards”—for using and viewing dashboards
- “Edit, save, and remove dashboards”—for editing, saving, and removing dashboards

The following steps outline the process of creating a dashboard:

- 1 Ensure that you have the privileges to create a dashboard.
- 2 Create a dashboard. See [“To add a dashboard:” on page 66](#).
- 3 Add panels to the dashboard you created. See [“To add a panel to a dashboard:” on page 68](#).

If you are adding a Search Results panel, the saved search must exist. If no saved searches exist, the Search Results panel option is not displayed.

To add a dashboard:

- 1 Click **Dashboards** from the top-level menu bar.

- 2 Click the **Tools** drop-down menu and select **Create Dashboard**.
- 3 Enter a meaningful name for the dashboard in the Name field.
- 4 Select whether the dashboard Type is Private or Shared.

The private dashboards are only visible to the user who created them, and the shared dashboards are visible to all users of Logger.

- 5 Click **Create**.

The dashboard is created. You must add panels to the dashboard next, as described in ["To add a panel to a dashboard:" on page 68](#).

To edit a dashboard:

When you edit a dashboard, you can change its name or privacy setting—Private or Shared. When you make a dashboard Shared, all Logger users can see it; however, they will not see the information to which they do not have privileges. For example, if a user does not have privileges to a storage group and a panel in a Shared dashboard includes a query that accesses the events in that storage group, the panel will be blank when the user accesses the shared dashboard.

- 1 Click **Dashboards** from the top-level menu bar.
- 2 Click the **Tools** drop-down menu and select **Edit Dashboard**.
- 3 If you want to change the name of the dashboard, enter a new name in the Name field.
- 4 If you want to change the privacy setting of the dashboard, select the appropriate setting from the Type drop-down menu, and click **Save**.

To delete a dashboard:

- 1 Click **Dashboards** from the top-level menu bar.
- 2 Select the dashboard that you want to delete.
- 3 Click the **Tools** drop-down menu and select **Delete Dashboard**.
- 4 Click **Yes** to confirm your action in the confirmation message, or click **No** to exit without making a change.

To set a dashboard as default:

When you set a dashboard as default, it is the default dashboard screen that displays when you navigate to the Dashboards menu. This setting is user-specific; therefore, your default dashboard can be different from that of another user.

For all Logger users, the Summary page (accessible from the Summary navigation option in the top-level menu bar) is the default home page. That is, when you log in, the Summary page is displayed first. However, you can configure Logger to display a specific dashboard as the default home page for you. To do so, first configure the option described in ["Options" on page 60](#), then

To select the dashboard you want to display by default:

- 1 Click **Dashboards** from the top-level menu bar.
- 2 Select the dashboard that you want to configure as default.
- 3 Click the **Tools** drop-down menu and select **Select as Default**.

- 4 Click **Yes** to confirm your action in the confirmation message, or click **No** to exit without making a change.

Adding and Managing Panels in a Dashboard

A dashboard can contain a mix of Search Results, Monitor, and Summary panels. There is no limit on the number of Monitor and Summary panels you add to a single dashboard; however, you can only add up to four Search Results panels for optimum performance.

You can add the following types of panels:

- Search Results—Chart and Table
- Monitor—All four types available under the default Monitor dashboard
- Summary—All four types available under the default Summary dashboard and user-defined Summary panels.

To add a panel to a dashboard:

- 1 Click **Dashboards** from the top-level menu bar.
- 2 Select the dashboard to which you want to add the panel.
- 3 Click the **Tools** drop-down menu and select **Add Panel**.
- 4 Configure these parameters and click **Add**.

Parameter	Description
Type	<p>Select the type of panel:</p> <ul style="list-style-type: none">• Search Result Chart—Displays search results in a chart form• Search Result Table—Displays search results in a table form• Monitor Graph—Displays a graph of the selected resource• Monitor Table (Forwarder)—Displays forwarder information in a table form• Monitor Table (Receiver)—Displays receiver information in a table form• Monitor Table (StorageGroup)—Displays storage group information in a table form• Summary (Agent Severities)—Displays event summary categorized by agent severities configured on your Logger• Summary (Agent Types)—Displays event summary categorized by receivers configured on your Logger• Summary (Receivers)—Displays event summary categorized by receivers configured on your Logger• Summary (Devices)—Displays event summary categorized by devices configured on your Logger• Summary (User Defined)—Displays event summary categorized by the field you select when adding the panel <p>Note: If no saved search queries exist on your Logger, the “Saved Search” panel types are not available as selections in the drop-down menu.</p>
Title	<p>Enter a meaningful name for the panel.</p> <p>A default name is present in this field, but you can change it.</p>

Parameter	Description
Graph	Only applicable to Monitor Graph panels. Select the type of graph you want the panel to display. Some of the available options are CPU Usage - 4 hour, Platform Memory Usage - Daily, and Disk Read-Write - Weekly.
Saved Search	<i>Only applicable to Search panels.</i> Select the saved search query to use for searching events that will be displayed in the panel.
Chart Type	<i>Only applicable to Search Result Chart panels.</i> Type of chart to display matching events. You can select from: Column, Bar, Pie, Area, Line, Stacked Column, Stacked Bar. Default: Column
Chart Limit	<i>Only applicable to Search Result Chart panels.</i> Number of unique values to plot. Default: 10
Field Name	Only applicable to Summary (User Defined) panels. The event field name by which the event summary on a Summary panel will be categorized. Default: agentSeverity

To edit a panel:

Once you add a panel to a dashboard, whether you can edit it depends on the type of panel. You can edit the Search Results panels and the user-defined Summary panels; the Monitor panels and some of the Summary panels are not editable.

The following table lists the panels you can edit and what you can edit in them.

Action	Description
All Panels	
Delete	Removes a panel from a dashboard.
Search Result Panels	
Edit Panel	Change Title, associated saved search, Chart Type, or Chart Limit
Edit Saved Search	Access the Edit Saved search page to edit the associated saved search query
View on Search Page	Runs the panel's query on the Search Results page (Analyze > Search) and displays matching events on that page
Refresh	Refreshes the current contents of the panel. Note: All other panel types are automatically refreshed; therefore, an explicit refresh is not required for them.
Summary Panels - User Defined	
Edit Panel	Change Title or field name by which events are categorized.

- 1 Click **Dashboards** from the top-level menu bar.
- 2 Select the dashboard that contains the panel you want to edit.
- 3 If you are editing a user-defined Summary panel:

- a** Click the (🔍) icon.
 - b** Edit the title, field name, or both.
- 4** If you are editing a Search Result panel:
 - a** Click the (📄) icon.
 - b** Select **Edit Panel** if you want to edit the panel title, select a different saved search; or, if applicable, chart type or chart limit.
 - c** Select **Edit Saved Search** if you want to access the Edit Saved Search page (under the Configuration menu option from the top-level menu bar) to edit the saved search query.
- 5** Click **Save**.

To delete a panel from a dashboard:

You cannot delete panels from the default Monitor dashboard or the default Summary dashboard that you access from the Monitor and Summary menu options from the top-level menu bar. However, Monitor and Summary panels added to the dashboards you created under the Dashboards menu option can be deleted.

- 1** Click **Dashboards** from the top-level menu bar.
- 2** Select the dashboard that contains the panel you want to delete.
- 3** Click the (✕) icon.
- 4** Click **Yes** to confirm your action in the confirmation message, or click **No** to exit without making a change.

To change the layout of a dashboard:

You can only change the layout of the dashboards you create. The Monitor dashboard layout cannot be changed.

- 1** Click **Dashboards** from the top-level menu bar.
- 2** Select the dashboard that contains the panel you want to rearrange.
- 3** Click the **Tools** drop-down menu and select **Change Layout**.
- 4** Point your cursor in the blue band that shows the panel title and drag the panel to a different position.
- 5** Click **Save** after you rearrange the panels.

The Default Monitor Dashboard

The Monitor Dashboard is available on your Logger by default. It displays the real-time and historical status of receivers, forwarders, and storage, CPU, and disk usage statistics. (On the software version of Logger, the CPU and disk usage statistics indicate the total use of these resources on the system, not just the use of these resources by the Logger process.)

The Summary panel, which is the default panel, shows the status of CPU Usage, Event Flow, Receivers, Forwarders, and Storage Groups in a summarized view. The other Monitor panels available through a drop-down menu are Platform, Network, Logger, Receivers, Forwarders, and Storage. You cannot change or adjust any panels available in this dashboard.

All monitor pages, except Summary, include an additional drop-down menu for duration control. On these, choose a time span for historical data:

- 4 hours
- 24 hours
- 7 days
- 30 days
- 90 days
- 365 days

On the Summary page, click on a Receiver, Forwarder, or Storage Group name to jump to the Search page and include the selected resource in the query.

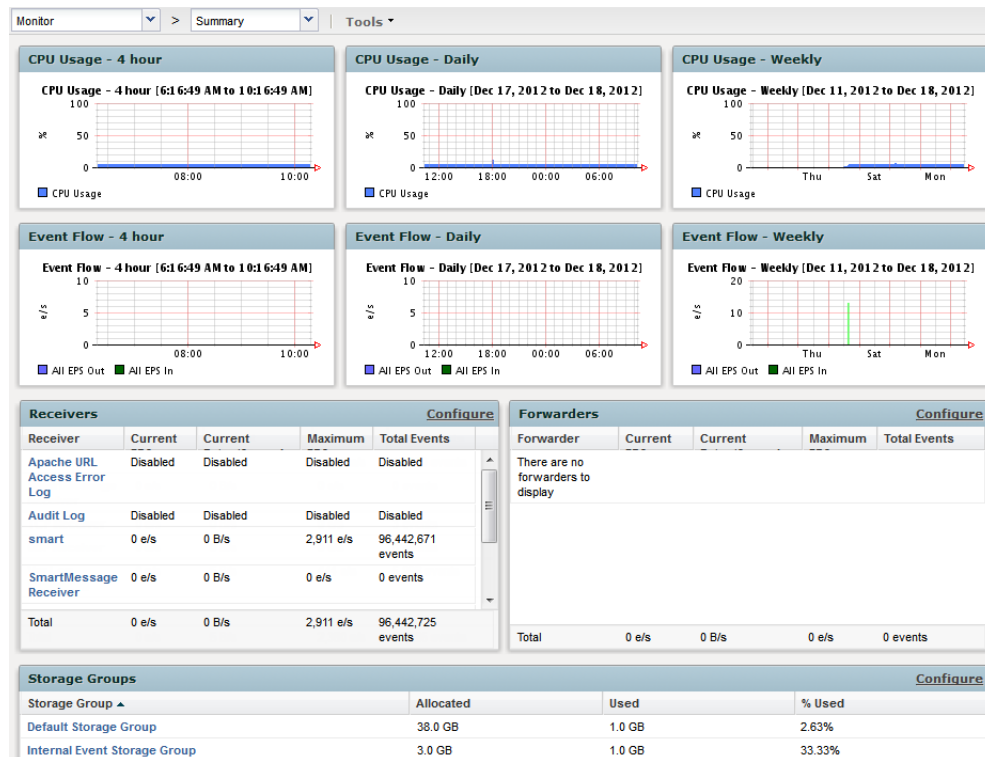


Figure 3-2 The Monitor dashboard displays summary information by default

The total space allocated for a storage group includes a certain amount that has been set aside to ensure that the group can receive new events when it is almost full. As a result, the percentage of used space for a storage group never reaches 100% (as displayed on the **Monitor > Summary** page). For software Loggers installed using the Minimal setting, the maximum % Used (On the **Monitor > Summary** page) for each storage group reaches up to 66.33%. (Two storage groups of 3 GB each; 1 GB is set aside for new events in each group. After 2 GB of space has been used and the new events are being written to the last 1 GB, Logger automatically triggers retention and reclaims 1 GB of the used space. Thus, the % Used field for each storage group only reaches up to 66.33%.)

The “Session Inactivity Timeout” setting on the Authentication Settings page (**System Admin > Users/Groups > Authentication**) does not apply to the user interface pages accessed through the Monitor menu. That is, if a user is on any of the user interface pages accessed through the Monitor menu and the session has been inactive for the number of

minutes specified in the “Session Inactivity Timeout” setting, the user’s session will not time out.

Platform

The Platform monitor page, as shown in [Figure 3-3](#), displays information about CPU usage, memory usage, bytes received and sent on the network, and raw disk reads and writes.




Figure 3-3 Monitor dashboard - Platform page

Network

The Network monitor page displays a graph for each network interface card. (The number of network interface cards varies by the hardware model.) The graph displays the bytes transmitted, overlaid on the bytes received.

Logger

The Logger monitor page, as shown in [Figure 3-4](#), displays information about events, searches, and memory. JVM Memory Usage chart displays the memory used by the Logger's back-end server process. For example, this could be the memory used to perform the search after receiving the search query from the UI.




Figure 3-4 Memory usage displayed on the Logger page of the Monitor dashboard

Receivers

The Receivers monitor page shows total Events per Second (EPS) received and displays values for each configured receiver. The list of receivers includes all receivers known to the system, including those that are disabled. To create a new receiver, or to enable or disable one, see [“Receivers” on page 260](#).

Forwarders

The Forwarders monitor page shows total Events per Second (EPS) sent and displays values for each configured forwarder. The list of forwarders includes all forwarders known to the system, including those that are disabled. To create a new forwarder, or to enable or disable one, see [“Forwarders” on page 285](#).

Storage

The Storage monitor page, shown in [Figure 3-5](#), displays disk read and disk write information. The list of storage groups compares allocated and used space in each group. Space is used in 1 GB chunks so a 5 GB storage group appears 20% used as soon as it is set up.

For more information about storage groups, see [“Storage Groups” on page 257](#).




Figure 3-5 Monitor dashboard - Storage page

Searching and Analyzing Events

This chapter describes how to search for specific events in Logger. First, the chapter discusses the methods available for search, how to query for events, how to save a defined query and the events that the query finds for future use. Next, the chapter describes how to set up alerts to be notified when events matching the criteria you specified are received.

[“The Need to Search Events” on page 75](#)
[“The Process of Searching Events” on page 76](#)
[“Elements of a Search Query” on page 78](#)
[“Syntax Reference for Query Expressions” on page 93](#)
[“Using the Advanced Search Builder Tool” on page 96](#)
[“Search Analyzer” on page 100](#)
[“Regex Helper Tool” on page 101](#)
[“Search Helper” on page 103](#)
[“Searching for Events on Logger” on page 106](#)
[“Understanding the Search Results Display” on page 109](#)
[“Exporting Search Results” on page 119](#)
[“Indexing” on page 121](#)
[“Saving Queries \(Saved Filters and Searches\)” on page 126](#)
[“System Filters/Predefined Filters” on page 128](#)
[“Alerts” on page 132](#)
[“Live Event Viewer” on page 133](#)

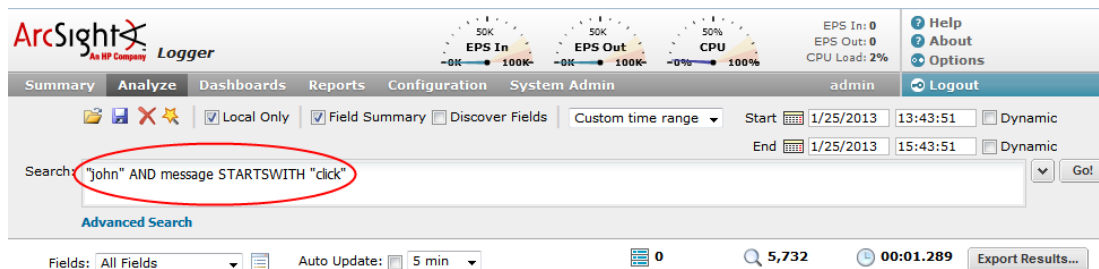
The Need to Search Events

When you want to analyze events matching specific criteria, include them in a report, or forward them to another system such as ArcSight ESM, you need to search for them. To search for events, you create queries. The queries you create can vary in complexity based on your needs. Queries can be simple search terms or they can be complex enough to match events that include multiple IP addresses or ports, and that occurred between specific time ranges from a specific storage group.

The Process of Searching Events

The search process uses an optimized search language that allows you to specify multiple search commands in a pipeline format. In addition, you can customize the display of search results, view search results as charts, and so on.

The most straightforward way to run a search is to enter the keywords or information you are searching for (the query) in the Search text box, select the time range, and click **Go!** You can enter a simple keyword, such as `hostA.companyxyz.com` or a complex query that includes Boolean expressions, keywords, fields, and regular expressions. The system searches for data that matches the criteria you specified and displays the results on the page where you entered your query.



The search results are displayed in a table and as a histogram as soon as they are returned, even if the query has not finished scanning all data. For an example, see [“Simple Query Example” on page 76](#).

You can also add a chart to your search to display the most important information in a more meaningful fashion. Charts are not displayed until all the data is returned. For an example, see [“Query Example Using a Chart” on page 77](#).

There are several convenient ways to enter a search query—Typing the query in the Search text box, using the Search Builder tool to create a query, or using a previously saved query (referred to as a filter or saved search).

When you type a query, the Search Helper provides suggestions and possible matches to help you build the query expression. (See [“Search Helper” on page 103](#) for more information.)

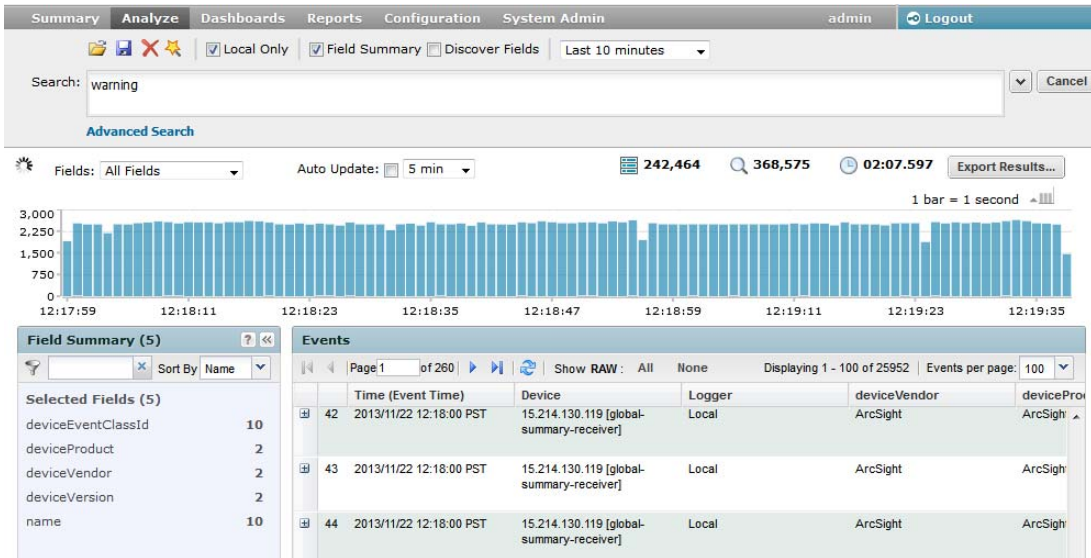
In addition to typing the query in the Search text box, you can do the following:

- Create queries by using the Advanced Search tool. For more information, see [“Using the Advanced Search Builder Tool” on page 96](#).
- Save queries and use them later. For more information, see [“Saving Queries \(Saved Filters and Searches\)” on page 126](#).
- Create new queries from the predefined queries that come with your system. For more information, see [“System Filters/Predefined Filters” on page 128](#).

Although a search query can be as simple as a keyword, you will be better able to utilize the full potential of the search operation if you are familiar with all the elements of a query, as described in [“Elements of a Search Query” on page 78](#).

Simple Query Example

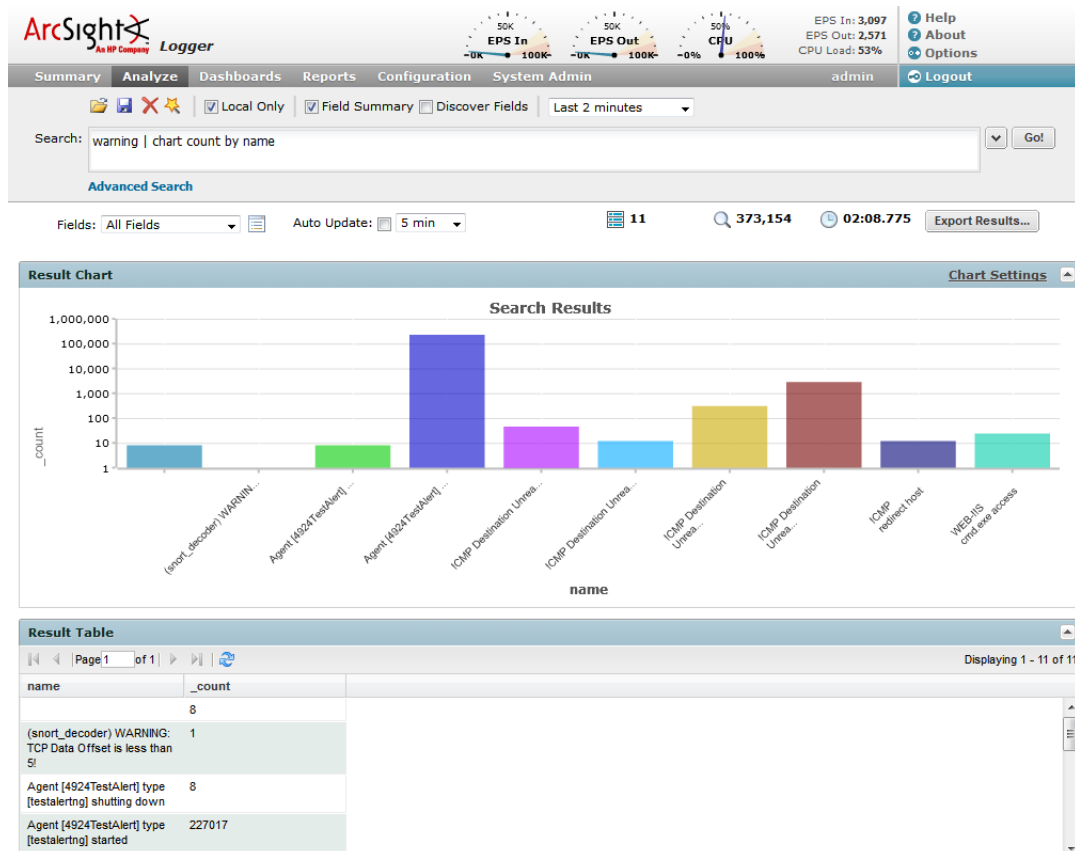
This example query finds events containing the word “warning”. Type `warning` in the search box and then click **Go!**



Query Example Using a Chart

Aggregated search operators such as chart, top, and rare generate charts of search results. This example query finds events containing the word “warning” and charts the number of warnings for each name. Type the following query in the search box and then click **Go!**

warning | chart count by name



For more information on the search operators, see [“Search Operators” on page 529](#). For more information on creating and using charts, see [“Chart Drill Down” on page 114](#) and [“Refining and Charting a Search from Field Summary” on page 118](#).

Elements of a Search Query

A simple search query consists of these elements:

- Query expression
- Time range
- Fieldset

An advanced Logger search query can also include constraints that limit the search to specific device groups, storage groups, and peer Loggers. For information about storage groups, devices groups, and peers, see [“Storage” on page 256](#), [“Device Groups” on page 249](#), and [“Peer Loggers” on page 326](#).

Query Expressions

A query expression is a set of conditions that are used to select events when a search is performed. An expression can specify a very simple term to match such as “login” or an IP address; or it can be more complex enough to match events that include multiple IP addresses or ports, and that occurred between specific time ranges from a specific storage group.

Specify the query in the Search text box by using the following syntax:

```
<Indexed Search> | <Search Operators>
```

The query expression is evaluated from left to right in a pipeline fashion. First, events matching the specified Indexed Search portion of the query are found. The search operator after the first pipe (“|”) character is then applied to the matched events followed by the next search operator, and so on to further refine the search results.

The search results table and the histogram display the events that match the query as they are found. As additional events are matched, the search results table and the histogram are refreshed. Certain search operators such as head and tail, require a query to finish running before search results can be displayed.

- The indexed search section of the query is described in [“Indexed Search” on page 78](#).
- The search operators are described in [“Search Operators” on page 86](#).

Indexed Search

The Indexed Search section of the query uses fields to search for relevant data quickly and efficiently. You can use a search expression to specify keywords to search for in the event text or to search using field-based expressions in a Boolean format.

- [“Keyword Search \(Full-text Search\)” on page 79](#)
- [“Field-based Search” on page 80](#)
- [“Searching for Rare Field Values” on page 83](#)

Keyword Search (Full-text Search)

Keywords are simply the words you want to search for, such as failed, login, and so on. You can specify multiple keywords in one query expression by using Boolean operators (AND, OR, or NOT) between them. Boolean expressions can be nested; for example, (John OR Jane) AND Doe*. If you need to search for the literal occurrence of AND, OR, or NOT (in upper-, lower-, or mixed case), enclose them in double quotes (") so the search engine does not interpret them as operators. For example, "and", "Or", and so on.



Although the Boolean operators AND, OR, and NOT can be specified in upper-, lower-, or mixed case when used as an operator, HP recommends that you use uppercase for ease of reading the query.

Keep the following in mind when specifying keyword search expressions:

- Follow the requirements described in ["Syntax Reference for Query Expressions" on page 93](#).
- Keyword search is not case sensitive.
- Use Boolean operators (AND, OR, or NOT) to connect multiple keywords. If no Boolean operator is specified between two keywords, the AND operator is applied by default. Also, use the Boolean operators to connect keywords to fields you specify.
- Use double quotes (") to enclose a single word for an exact match. Otherwise, the word is treated as <search string>*. For example, to search for log, type "log". If you type log (without the double quotes), the search will match all words that begin with log; for example, log, logger, logging, and so on.
- When specifying Boolean operators (AND, OR, or NOT) as keywords, enclose them in double quotes ("). For example, "AND".
- Use the backslash (\) as an escape character for \, ", and *. However, backslash will not escape these characters if the keyword is enclosed in double quotes. For example, "log\ger" and log\ger will match the same values—log\ger in both cases. Likewise, log*ger and "log*ger" will match the same values—log*ger, in this case.

The following table summarizes how special characters are treated in a keyword search.

Character	Usage
Space	You cannot specify keywords that contain the characters in the left column. Therefore, to search for a phrase such as <i>failed login</i> , enter "failed" AND "login".
Tab	
Newline	Note: * is a valid character for wildcard character searches.
,	
;	
(
)	
[
]	
{	
}	
"	
*	

Character	Usage
=	<p>To specify a keyword that contains any of the characters in the left column, enclose the keyword in double quotes (" "). You can also specify an asterisk (*) at the end of the keyword for an exact match.</p> <p>Examples:</p> <ul style="list-style-type: none"> "C:\directory" "result=failed"
:	
/	
\	
@	
-	
?	
#	
\$	
&	
_	
%	
>	
<	
!	
*	<p>You can use the wildcard character asterisk (*) to search for keywords, however, the wildcard cannot be the leading character in the keyword. Therefore, the following usage is valid:</p> <ul style="list-style-type: none"> log* "log*" log* log* log*app log*app*app <p>However, the following usage is not valid:</p> <ul style="list-style-type: none"> *log *log*app*

Field-based Search

The Logger schema contains a predefined set of fields. You can add fields that are relevant to the events you collect on your Logger to its schema. A field-based search can only contain fields in Logger's schema. (See additional guidelines at ["Guidelines for field-based search expressions:" on page 82.](#))

The Logger indexing capability allows for schema fields to be indexed. Logger's search operation and reports utilize the indexed fields to yield significant search and reporting performance gains. Although you can add indexed and non-indexed fields to a search query, **search and reporting performance will be much faster if all fields in a query are indexed**. For more information and a list of fields you can index, see ["Indexing" on page 121](#). For discussion on field-based query performance, see ["Performance Optimizations for Indexed Fields in Search Queries" on page 101](#).

The field operators you can use in a query expression are listed in the table below. In addition to the field operators, you can use search operators, as discussed in ["Search Operators" on page 86](#).

You can specify multiple field conditions in one query expression by using the listed operators between them. The conditions can be nested; for example, (name="John Doe" OR name="Jane Doe") AND message!="success".



Note

If a query includes the Boolean operator OR and the metadata identifiers (discussed in [“Constraints” on page 92](#)), the expression to be evaluated with OR must be enclosed in parentheses, as shown in this example:

```
(success OR fail) _storageGroup IN ["Default Storage Group"]
```

If the expression is not enclosed in parentheses, an error message is displayed.

Any literal operator in the table can be specified in upper-, lower-, or mixed case. To search for these words as literals in events, enclose them in double quotes (""). For example, message CONTAINS "Between".

To determine the data type of a field, see [“Viewing Default Fields” on page 324](#).

To determine the size of a custom field, see [“Viewing Custom Fields” on page 324](#).

Table 4-1 Operators for field based search

Operator	Example	Notes
AND	name="Data List" AND message="Hello" AND 1.2.3.4	Valid for all data types.
OR	(name="TestEvent" OR message="Hello") AND type=2 AND 1.2.4.3	Valid for all data types.
NOT	NOT name="test 123"	Valid for all data types.
!=	destinationPort != 100 message!="failed login" message!=failed*login (* means wildcard) "test" message!=failed*login (* is literal in this case)	Valid for all data types.
=	bytesIn = 32 message="failed login" message="failed*login" (* means wildcard)	Valid for all data types. The size of each field in the schema is predetermined. If the string you are searching for is longer than the field-length, you should use a STARTSWITH rather than an = search, and include no more than the number of characters in the field size. To determine the size of a default field, see “Viewing Default Fields” on page 324 . To determine the size of a custom field, see “Viewing Custom Fields” on page 324 .

Table 4-1 Operators for field based search (Continued)

Operator	Example	Notes
> *	bytesIn > 100	Valid for all data types.
< *	startTime < "\$Now - 1d"	
>= *	endTime >= "01/13/2014 07:07:21" endTime >= "2014/13/01 00:00:00 PDT" endTime >= "Sep 10 2014 00:00:00 PDT"	* These operators evaluate the condition lexicographically. For example, deviceHostName BETWEEN AM AND EU searches for all devices whose names start with AM, AMA, AMB, AN, AO, AP and so on, up to EU. Therefore, any device whose name starts with AK, AL, and so on is ignored. Similarly, devices with names EUA, EUB, FA, GB, and so on will be ignored.
<= *	startTime <= "\$Now - 1d"	
IN*	priority IN [2,5,4,3] destinationAddress IN ["10.0.20.40", "209.128.98.147"] _deviceGroup IN ["DM1"] _storageGroup NOT IN ["Internal Event Storage Group", "SG1"] _peerLogger IN ["192.0.2.10", "192.0.2.11"]	
BETWEEN*	priority BETWEEN 1 AND 5	
STARTSWITH	message STARTSWITH "failed"	Valid for string (text) data types only.
ENDSWITH	message ENDSWITH "login"	Valid for string (text) data types only.
CONTAINS	message CONTAINS "foobar"	Valid for string (text) data types only.
IS	sessionId IS NULL sessionId IS NOT NULL	Valid for all data types.

Follow these guidelines when specifying field-based search expressions.

Guidelines for field-based search expressions:

- Follow the requirements described in ["Syntax Reference for Query Expressions" on page 93](#).
- For faster searches, follow the recommendations in ["Searching for Rare Field Values" on page 83](#) and ["Tuning Search Performance" on page 108](#).
- By default, field-based search is case sensitive. You can change the sensitivity from the Field Search Options section of the **Configuration** (or **Configuration > Settings**) **Administration > Search > Search Options** tab. For more information, see ["Tuning Advanced Search Options" on page 321](#).
- You can specify any predefined Logger schema field. For example, cat = /Monitor/CPU/Usage. For a complete list, see ["Indexing" on page 121](#).
- You can specify any custom field you have added to the schema. For example, SSN=333-333-3333. For more information about custom schema fields, see ["Adding or Importing Schema Fields" on page 345](#).
- You cannot specify user-defined fields created through a predefined or user-defined parser in the Indexed Search portion of a query. (The Indexed Search portion of a query is the expression before the first pipeline character.)
- A query expression (Indexed Search | Search Operators) is evaluated from left to right in a pipeline fashion. By design, a parser—predefined or user-defined—is applied to an event when the Search Operators are processed in a search query. Therefore, field

creation when a parser is applied to an event occurs later than the Indexed Search stage. As a result, you cannot specify these fields in a field-based search query.

For example, the Apache Access Log parser creates the field SourceHost. You cannot specify the following query expression:

```
SourceHost="192.0.2.0"
```

However, you can use this field after the first pipeline, as shown in this example.

```
| where SourceHost="192.0.2.0"
```

Or, if you want to search only the Apache Access Logs for

SourceHost="192.0.2.0", you can specify this expression:

```
| where parser="Apache Access Log" and clientIP="192.0.2.0"
```

Additionally, you can run a full-text (keyword) search on "192.0.2.0", as follows:

```
"123.456.789" | where SourceHost="192.0.2.0"
```

- If an event field contains data of an unexpected type (for example, a string when an integer is expected), the data is ignored. Therefore, search for that data value will not yield any results. For example, if the port field contains a value 8080A (alphanumeric) instead of 8080 (numeric), the alphanumeric value is ignored. The datatypes of the schema fields are available from the **Configuration** (or **Configuration > Settings**) **Administration > Search > Default Fields** tab. For more information on how to view this information, see ["Viewing Default Fields" on page 324](#).
- For optimal search performance, make sure that event fields on ALL peers are indexed for the time range specified in a query. If an event field is indexed on one system but not on its peers for a specific time range, a distributed search will run slower on the peers. However, it will run at optimal speed on the local system. Therefore, the search performance in such a setup will be slow.
- For faster report generation, ALL fields of a report (including the fields being displayed in the report) need to be indexed. That is, in addition to the fields in the WHERE clause of the query, the fields in the SELECT clause also need to be indexed.

Searching for Rare Field Values

To enable you to quickly search common IP address, host name, and user name fields for rare values; Logger 5.5 and later creates super indexes on new data as it comes in. For more information, see ["Super Indexing" on page 125](#).

Searches written to take advantage of super indexes will tell you very quickly if there are no hits and will return results more quickly than regular searches when there are very few hits. Thus, they are excellent for fast needle-in-a-haystack searches.



Since super indexes are built on new data as it comes in on, they only apply to data collected by Logger 5.5 or later. Any data brought forward from an upgrade from an earlier version of Logger will not be super-indexed and thus will not exhibit this search speed improvement.

Writing Searches to Increase Search Speed on Super-Indexed Fields

To take advantage of super indexing and get the fastest search results, run an equal to (=) search, such as `sourceAddress=192.0.2.0`, and write the indexed search portion of

your query to include uncommon values in the super-indexed fields listed in the table below.

Table 4-2 Fields With Super Indexes

deviceEventClassId	deviceProduct	deviceVendor	destinationHostName
destinationPort	destinationAddress	destinationUserId	destinationUserName
deviceAddress	deviceHostName	sourceHostName	sourcePort
sourceAddress	sourceUserId	sourceUserName	



Unlike the indexed fields discussed in [“Field-based Indexing” on page 122](#), you cannot add to the list of super-indexed fields.

Super indexes speed up searches that use the equal to (=) operator in the indexed search portion of the query expression. They have no performance impact on searches that use greater than (>), less than (<), not equal to (!=), or other operators in the indexed search portion of the query. While Logger supports full-text search, search on fields that are not super-indexed, and searches that use operators such as >, less than <, !=, and so on; such searches may not provide the greatest search speed.

Using AND and OR with the = operator can be very powerful when searching super-indexed fields. However, to obtain the greatest search speed improvement, you must use them carefully. The table below provides examples to help you understand how to write queries that take advantage of the power of super indexing.

Table 4-3 Query Examples for Super-Indexed Needle-in-a-Haystack Searches

Query	Does It Improve Search Speed?
Arcsight (full text)	No difference. This is a full text query, and so does not take advantage of super-indexed field search speed improvements.
192.0.2.0 (full text that looks like a super-indexed field)	No difference. While this could be an IP address, it is a full text search, not an = search against one of the super-indexed fields, and so does not take advantage of super-indexed field search speed improvements.
sourceAddress = 192.0.2.0 (= on a super-indexed field)	The search speed is improved and the results return very quickly when there are no hits. If Logger has not encountered 192.0.2.0 as a sourceAddress, it quickly returns the message "No results were found". If it has encountered that sourceAddress, the range of events to be searched is narrowed down.
sourceAddress = 192.0.2.0 OR sourceAddress = 192.0.2.2 (= using OR on super-indexed fields)	The search speed is improved and the results return very quickly when there are no hits. If Logger has not encountered 192.0.2.0 or 192.0.2.2 as a sourceAddress, it quickly returns the message "No results were found". If it has encountered one or the other, the range of events to be searched is narrowed down.

Table 4-3 Query Examples for Super-Indexed Needle-in-a-Haystack Searches (Continued)

Query	Does It Improve Search Speed?
sourceAddress = 192.0.2.0 AND destinationAddress = 192.0.2.2 (= using AND on super-indexed fields)	<p>The search speed is improved and the results return very quickly when there are no hits.</p> <p>If Logger has not encountered 192.0.2.0 as a sourceAddress, it quickly returns the message "No results were found".</p> <p>Similarly, if Logger has not encountered 192.0.2.2 as a destinationAddress, it quickly returns the message "No results were found", even if it has encountered 192.0.2.0 as a sourceAddress.</p> <p>If Logger has encountered both, the range of events to be searched is narrowed down.</p>
sourceAddress != 192.0.2.0 (!= on a super-indexed field)	<p>No difference.</p> <p>Super indexing does not help with negations, so this query does not take advantage of super-indexed field search speed improvements.</p>
sourceAddress != 192.0.2.0 OR destinationAddress= 192.0.2.2 (!= using OR on Super-indexed fields)	<p>No difference.</p> <p>Since there is a negation on the sourceAddress and this is an OR condition, this query does not take advantage of super-indexed field search speed improvements.</p>
sourceAddress != 192.0.2.0 AND destinationAddress = 192.0.2.2 (!= using AND on Super-indexed fields)	<p>The search speed is improved and the results return very quickly when there are no hits.</p> <p>Since this is an AND condition, both conditions need to be true.</p> <p>Even though there is a negation on the sourceAddress, if Logger has not encountered a destinationAddress address of 192.0.2.2, this AND condition will never be satisfied. In that case, it quickly returns the message "No results were found".</p> <p>If Logger has encountered that destinationAddress, the range of events to be searched is narrowed down.</p>
sourceAddress = 192.0.2.0 AND arcsight (= on super-indexed field AND full text)	<p>The search speed is improved and the results return very quickly when there are no hits.</p> <p>If Logger has not encountered a sourceAddress of 192.0.2.0, this AND condition will never be satisfied. In that case, it quickly returns the message "No results were found", even though there is a full text search.</p> <p>If Logger has encountered that sourceAddress, the range of events to be searched is narrowed down.</p>
sourceAddress = 192.0.2.0 OR arcsight (= on super-indexed field OR full text)	<p>No difference.</p> <p>Regardless of whether Logger has encountered a sourceAddress of 192.0.2.0, the OR condition requires a full text search for "arcsight", so this query does not take advantage of super-indexed field search speed improvements.</p>
requestMethod = GET AND sourceAddress = 192.0.2.0 (NON-super-indexed field AND super-indexed field)	<p>The search speed is improved and the results return very quickly when there are no hits.</p> <p>Even though requestMethod is not one of the super-indexed fields, because the query uses an AND condition, Logger quickly returns the message "No results were found" if it has not encountered a sourceAddress of 192.0.2.0.</p> <p>If Logger has encountered that sourceAddress, the range of events to be searched is narrowed down.</p>

Table 4-3 Query Examples for Super-Indexed Needle-in-a-Haystack Searches (Continued)

Query	Does It Improve Search Speed?
requestMethod = GET OR sourceAddress = 192.0.2.0 (NON-super-indexed field OR super-indexed field)	No difference. Even though sourceAddress is one of the super-indexed fields, because it is in an OR condition with requestMethod, which is not super-indexed, this query does not take advantage of super-indexed field search speed improvements.
sourceAddress = 192.0.2.0 AND (sourceHostName = myhost.com OR sourcePort = 80) AND (destinationAddress = 192.0.2.2 OR arcsight) (super-indexed field AND (nested OR condition) AND (nested OR condition))	Results return very quickly when there are no hits. If Logger has not encountered a sourceAddress of 192.0.2.0, the top level AND will never be true. It quickly returns the message "No results were found" in that case. If Logger has not encountered a sourceHostName of myhost.com AND it has not encountered a sourcePort of 80, then the OR condition will never be true. Thus the top level AND condition will never be true. It quickly returns the message "No results were found" in that case. If Logger cannot show that the above conditions are false, then there will be no difference in search speed. Even though destinationAddress is one of the super-indexed fields, because it is in an OR condition with a full-text search for "arcsight", the range of events to be searched cannot be narrowed down.

Search Operators

The Search Operators portion of the query enables you to further refine the data that matched the indexed search filter. See [“Search Operators” on page 529](#) for a list of search operators and examples of how to use them.

The `rex` search operator is useful for syslog events (raw or unstructured data) or if you want to extract information from a specific point in an event, such as the 15th character in an event. Other operators such as `head`, `tail`, `top`, `rare`, `chart`, `sort`, `fields`, and `eval` are applied to the fields you specify or the information you extract using the `rex` operator.

Prior to Logger 5.2, you needed to use a special search operator—`cef`—to extract CEF fields from CEF events (structured data) that matched the indexed search filter (the query portion before the first pipeline in the query expression) before you could use other search operators to act upon those fields. However, starting with Logger 5.2, you do not need to explicitly extract the CEF fields and then apply other search operators to those fields. You can specify the event fields directly in queries. The CEF operator has been deprecated as of Logger 5.2.

Time Range

An event is timestamped with the receipt time when it is received on the Logger. **A search query uses this time to search for matching events.** Under most circumstances, the Logger receipt time is same as the event time. However, the event time and the Logger receipt time for an event can be different because there is usually a small lag between the time an event leaves a device and it is received at the Logger. If the device's clock is ahead or behind the Logger clock, the lag or lead can be significant.

A search operation requires you to specify the time range within which events would be searched. You can select from many predefined time ranges or define a custom time range to suit your needs.

Predefined time range: When you select a predefined time range such as “Last 2 Hours” or “Today”, the time range is relative to the current time. For example, if you select “Last 2 Hours” at 2:00:00 p.m. on July 13th, events from 12:00:00 to 2:00:00 p.m. on July 13th will be searched. If you refresh your search results at 5:00:00 p.m. on the same day, the time window is recalculated. Therefore, events that match the specified criteria and occurred between 3:00:00 and 5:00:00 p.m. on July 13th are displayed.


Custom time range: You can specify a time range in a 24-hour format to suit your needs. For example, a custom time range is:

Start: 8/13/2014 13:36:30
End: 8/13/2014 22:36:30

By default, the end time for a custom time range is the current time on your Logger and the start time is two hours before the current time.

You can also use variables to specify custom time ranges. For example, a dynamic date range might start at \$Now - 2h (two hours ago) and end at \$Now (the current time). The dynamic search is relative to when the query is run. Scheduled search operations use this mechanism to search through newer event data each time they are run.

The “Dynamic” field in the user interface enables you to specify the dynamic time, as shown in the following figure:



Following is a typical example of a dynamic search that limits results to the last two hours of activity:

Start: \$Now - 2h
End: \$Now

The syntax for dynamic search is:

<current_period> [+/- <units>]

Where <current_period>, such as \$Now, either stands alone or is followed by either a plus ('+') or minus ('-') and a number of units, such as 2h for two hours. The <current_period> always starts with a '\$' and consists of a word, case-sensitive, with no spaces, as shown in

Table 4-4 on page 88. The <units> portion, if given, consists of an integer and a single, case-sensitive letter, as shown in Table 4-5 on page 88.

Table 4-4 Current Period

Period	Description
\$Now	The current minute
\$Today	Midnight (the beginning of the first minute) of the current day
\$CurrentWeek	Midnight of the previous Monday (or same as \$Today if today is Monday)
\$CurrentMonth	Midnight on the first day of the current month
\$CurrentYear	Midnight on the first day of the current year

Table 4-5 Units

Unit	Description
m (lowercase)	Minutes (Do not confuse with 'M', meaning months)
h	Hours
d	Days
w	Weeks
M (uppercase)	Months (Do not confuse with 'm', meaning minutes)


Fieldsets

A fieldset determines the fields that are displayed in the search results for each event that matched a search query. The system provides a number of predefined fieldsets.



Note

The first time you open the search page in a new browser window the fieldsets lists are hidden and you cannot select them. Run a short search to display the hidden options.


- To view the current list of available fieldsets, click the down arrow in the Fields dialog box. The current System Fieldsets list is displayed. To see the fields included in each of the predefined fieldsets, click the  (Customize Fieldset) icon. To view a list of fields that are included in the search results for each fieldset type, select the fieldset from the drop-down list and mouse over the Field's label.




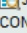

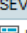
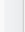
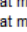
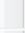
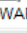
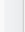
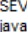








Note

Only fields available for matched events are displayed in a Search Results display (or the exported file). Therefore, even if you select the All Fields fieldset, you might not see all fields displayed in the search results.

- When you use a search operator that defines a new field, such as rex, rename, or eval, a new column for each field is added to the currently selected display. These newly defined fields are displayed by default. The fieldset, User Defined Fields, enables you to view only the newly defined fields.

- Prior to 5.2, you could only view the raw events if you clicked the  icon in Search Results; the raw event would be displayed in an unformatted text form. Starting with Logger 5.2, the predefined fieldset, Raw Event, is available. This fieldset displays the whole raw syslog event in a column called rawEvent, as shown in the following figure. The event is formatted to fit in the column.

Events			
<div> <div> <div>Page 1 of 149</div> <div> <div></div> <div></div> <div></div> </div> </div> <div>Show RAW : All None</div> </div>			
	Time (Event Time)	Device	rawEvent
	1 2011/10/21 11:32:20 PDT	10.4.13.110 [TCP Receiver]	 Jan 8, 2003 10:43:47 AM myapp.business.SimpleLogger doSomething SEVERE: Some message java.lang.IllegalArgumentException: Some exception text at myapp.business.SimpleLogger.doSomething(SimpleLogger.java:39) at myapp.business.SimpleLogger.main(SimpleLogger.java:13)
	2 2011/10/21 11:32:20 PDT	10.4.13.110 [TCP Receiver]	 Jan 8, 2003 10:43:46 AM myapp.business.SimpleLogger doSomething CONFIG: this is config
	3 2011/10/21 11:32:20 PDT	10.4.13.110 [TCP Receiver]	 Jan 8, 2003 10:43:47 AM myapp.business.SimpleLogger doSomething WARNING: this is a warning
	4 2011/10/21 11:32:20 PDT	10.4.13.110 [TCP Receiver]	 Jan 8, 2003 10:43:47 AM myapp.business.SimpleLogger doSomething SEVERE: this is severe
	5 2011/10/21 11:32:20 PDT	10.4.13.110 [TCP Receiver]	 Jan 8, 2003 10:43:47 AM myapp.business.SimpleLogger doSomething SEVERE: Some message java.lang.IllegalArgumentException: Some exception text at myapp.business.SimpleLogger.doSomething(SimpleLogger.java:39) at myapp.business.SimpleLogger.main(SimpleLogger.java:13)
	6 2011/10/21 11:32:20 PDT	10.4.13.110 [TCP Receiver]	 Jan 8, 2003 10:52:31 AM myapp.business.SimpleLogger doSomething INFO: this is info
	7 2011/10/21 11:32:20 PDT	10.4.13.110 [TCP Receiver]	 Jan 8, 2003 10:52:31 AM myapp.business.SimpleLogger doSomething WARNING: this is a warning
	8 2011/10/21 11:32:20 PDT	10.4.13.110 [TCP Receiver]	 Jan 8, 2003 10:52:31 AM myapp.business.SimpleLogger doSomething SEVERE: this is severe
	9 2011/10/21 11:32:20 PDT	10.4.13.110 [TCP Receiver]	 Jan 8, 2003 10:52:31 AM myapp.business.SimpleLogger doSomething SEVERE: Some message java.lang.IllegalArgumentException: Some exception text at myapp.business.SimpleLogger.doSomething(SimpleLogger.java:39)

Although the Raw Event field is most applicable for syslog events, you can also display the raw event associated with CEF events in the rawEvent column. To do so, make sure the connector that is sending events to the Logger populates the rawEvent field with the raw event.

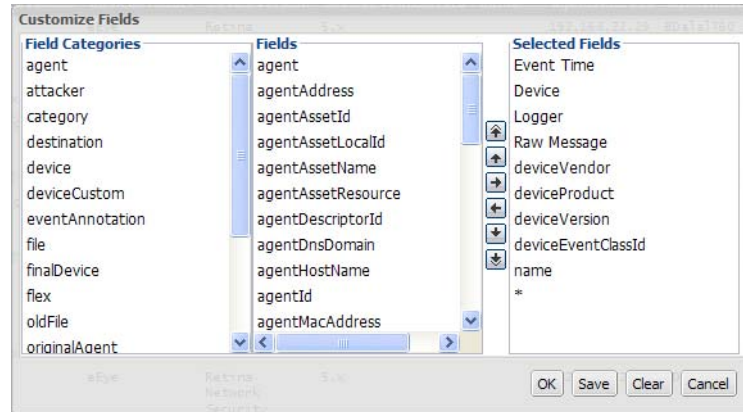


Note


To see the raw events in the rawEvent column, enable the Search Option, "Populate rawEvent field for syslog events". See ["Tuning Advanced Search Options" on page 321](#) for more information.

Creating Custom Fieldsets

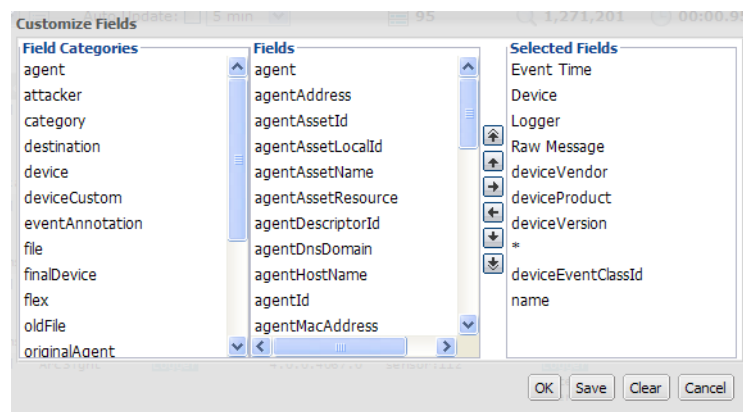
You can also create your own fieldsets by selecting "Customize..." from the "Fields" drop-down menu. The user interface offers a simple and intuitive way to select and move event fields you want to include in a fieldset, as shown in the following figure.



Use these buttons to create and edit a custom fieldset.



A wildcard field ("*") is available in the Fields list when you create a custom fieldset. This field includes all fields available in an event that are not individually listed in the custom fieldset definition. For example, for the following custom fieldset definition, the search results will list the fields before the asterisk ("*") first, followed by any other fields in an event. Lastly, the deviceEventClassId and Name fields will be listed.



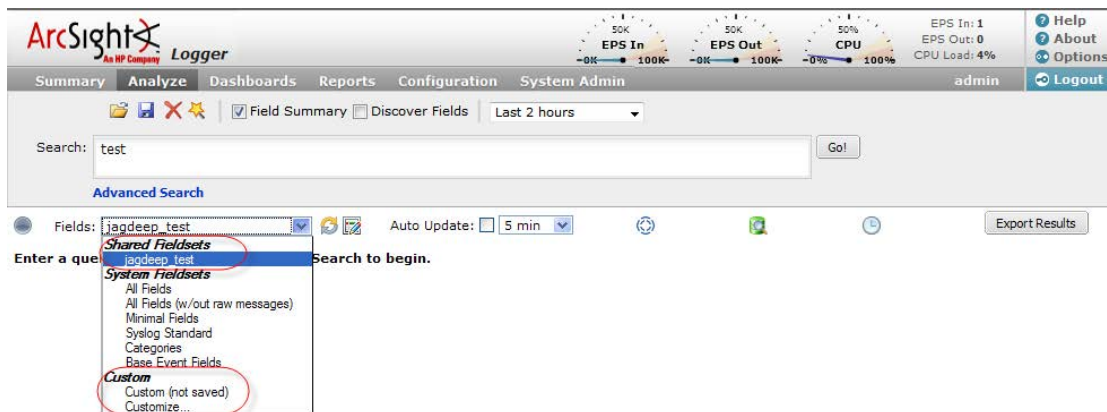
You can save the custom fieldset or use it only for the current session.



If you click **OK**, the fieldset appears in the Custom category. It is labeled as “Custom (not saved)” and is not visible to other users. It will remain available to you for this session. Once you log out of the current session, the temporary fieldset will be deleted. You can only have one temporary custom fieldset at a time.

If you click **Save**, the fieldset appears under the Shared Fieldsets category and is visible and available to the other users, as shown in the following figure. After a fieldset is saved, you can edit and delete it.

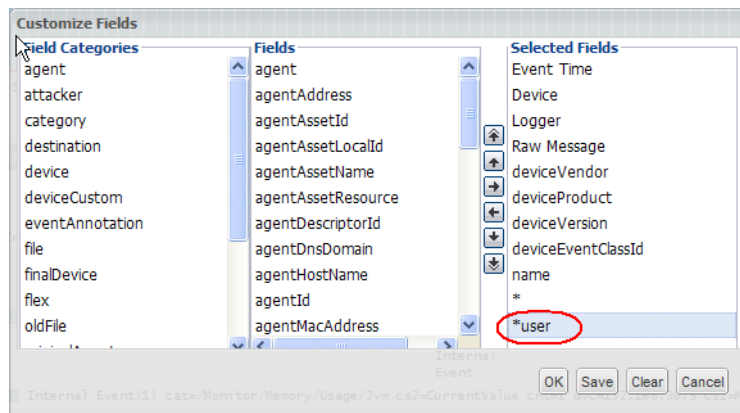
When saving a custom fieldset, you can specify it as the default for this system. If you do so, it is the default fieldset for all users on that system. If do not select it as the default, the fieldset is used only for your search results and does not affect other users connecting to the same system.



Fieldsets are not included in the saved filter definition.

For information about deleting custom fieldsets, see [“Viewing and Deleting Field Sets” on page 323](#)

The *user field, shown below, controls the display of fields defined by search operators (rex, rename, extract, or eval) as well as the fields created when a parser is applied to an event. When *user is included in the Selected Fields list of a custom fieldset, the created or defined fields are displayed.



Constraints

Using constraints in a query can speed up a search operation as they limit the scope of data that needs to be searched. Constraints enable you to limit a query to events from one or more of the following:

- Particular device groups
- Particular storage groups
- Specific peers

For example, you might want to search for events in the SG1 and SG2 storage groups on the local system only.

For information about storage groups and peers, see ["Storage" on page 256](#) and ["Device Groups" on page 249](#).

Follow these guidelines when specifying constraints:

- Use the following operators to specify constraints in a search query expression:

Metadata Identifier	Example
<code>_deviceGroup</code>	<code>_deviceGroup IN ["DM1", "HostA"]</code> where DM1 is a device group, while HostA is a device. Note: You can use this field to specify individual devices, as shown in the example above.
<code>_storageGroup</code>	<code>_storageGroup IN ["Internal Event Storage Group", "SG1"]</code>
<code>_peerLogger</code>	<code>_peerLogger IN ["192.0.2.10", "192.0.2.11"]</code>

- If a query includes the Boolean operator OR and metadata identifiers, the expression to be evaluated with OR must be enclosed in parentheses, as shown in this example:

```
(success OR fail) _storageGroup IN ["Default Storage Group"]
```

If the expression to be evaluated with OR is not enclosed in parentheses, an error message is displayed on the user interface screen.

- When specifying multiple groups in a constraint, ensure that the group names are enclosed in square brackets; for example, `_storageGroup IN ["SGA", "SGB"]`.
- You can apply constraints to a search query by:

- ◆ Typing the constraint in the Search text box

Once you type `"_s"` (for storage group), `"_d"` (for device group), or `"_p"` (for peer) in the Search text box, Search Helper automatically provides a drop-down list of relevant terms and operators from which you can select.



If a search query contains constraints and a regular expression, make sure that the constraints are specified before the regular expression. For example, `_peerLogger IN ["192.0.2.10"] name contains abc | REGEX=":\d31"`

- ◆ Selecting Storage Groups or peers from the Advanced Search tool. (To access the Advanced Search tool, click **Advanced Search** beneath the text box where you type the query.) For more information about the Advanced Search, see [“Using the Advanced Search Builder Tool” on page 96](#).

Syntax Reference for Query Expressions

To create valid and accurate query expressions, follow these requirements.

Table 4-6 Query Syntax Requirements

Behavior	Full Text Search	Field Search	Regular Expression
Case sensitivity	Insensitive (Cannot be changed.)	Sensitive (Can be changed using Tuning options. See “Tuning Advanced Search Options” on page 321 .)	Insensitive (Can be changed using Tuning options. See “Tuning Advanced Search Options” on page 321 .)
Escape character	\ Use to escape \. You cannot escape any other character.	\ Use to escape \, ", and *. Examples: name=log\\ger (matches log\ger) • name=logger* (matches logger*)	\ Use to escape any special character. Example: To search for a term with the character "[": REGEX= "logger\[
Escaping wildcard character	Cannot search for * Example: log* is invalid	Can search for * by escaping the character name=log* is valid	Can search for * by escaping the character
Exact Match/Search string includes an operator or a special character	Enclose keyword in double quotes; Otherwise, keyword treated as keyword*. Example: log (matches log, logging, logger, and so on) "log" (matches only log) Note: See the list of special characters that cannot be searched even when enclosed in double quotes, later in this table.	Enclose value in double quotes Example: message="failed login"	No special requirement.

Table 4-6 Query Syntax Requirements (Continued)

Behavior	Full Text Search	Field Search	Regular Expression
Nesting (including parenthetical clauses, such as (a OR b) AND c)	<p>Allowed</p> <ul style="list-style-type: none"> Use Boolean operators to connect and nest keywords. Metadata identifiers (<code>_storageGroup</code>, <code>_deviceGroup</code>, and <code>_peerLogger</code>), but can only appear at the top level in a query expression). If the query contains a regular expression, the metadata identifiers need to precede the regular expression. 	<p>Allowed</p> <ul style="list-style-type: none"> Use any operator listed in the “Field-based Search” on page 80 section to connect and nest field search expressions. Metadata identifiers (<code>_storageGroup</code>, <code>_deviceGroup</code>, and <code>_peerLogger</code>), but can only appear at the top level in a query expression 	<p>Multiple regular expressions can be specified in one query using this syntax:</p> <pre> REGEX= "<REGEX1>" REGEX= "<REGEX2>" ...</pre>
Operators	<p>Upper-, lower-, or mixed case Boolean operators—AND, OR, NOT. If an operator is not specified, AND is used.</p> <p>To search for literal operator AND, OR, NOT, in an event, enclose them in double quotes.</p> <p>Example: “AND”, “or”, “Not”</p> <p>Note: If a query includes the Boolean operator OR and the metadata identifiers (<code>_storageGroup</code>, <code>_deviceGroup</code>, and <code>_peerLogger</code>), the expression to be evaluated with OR must be enclosed in parentheses, as shown in this example:</p> <pre>(success OR fail) _storageGroup IN ["Default Storage Group"]</pre>	<p>Use any operator listed in the “Field-based Search” on page 80 section.</p> <ul style="list-style-type: none"> Unless a value is enclosed between double quotes, a space between values is interpreted as an AND. For example, <code>name=John Doe</code> is interpreted as <code>John AND Doe</code>. If an operator is not specified between multiple field expressions, AND is used. To search for literal operator, enclose the operator in double quotes. Examples: <pre>message STARTSWITH="NOT" message="LOGIN DID NOT SUCCEED"</pre> If a query includes the Boolean operator OR and the metadata identifiers (<code>_storageGroup</code>, <code>_deviceGroup</code>, and <code>_peerLogger</code>), the expression to be evaluated with OR must be enclosed in parentheses, as shown in this example: <pre>(success OR fail) _storageGroup IN ["Default Storage Group"]</pre> 	<p> and the operators described in “Time Range” on page 86.</p> <p>Use this operator to AND multiple regular expressions in one query expression.</p>

Table 4-6 Query Syntax Requirements (Continued)

Behavior	Full Text Search	Field Search	Regular Expression
Primary Delimiters: Space ' ; () [] } " * > < !	You can search for keywords containing primary delimiters by enclosing the keywords in double quotes. Example: "John Doe" "Name=John Doe" "www.hp.com"	You can search for these characters. Enclose value in double quotes if value contains any of these characters. Example: name="John*"	<ul style="list-style-type: none"> Cannot contain ^ in the beginning and \$ at the end as a matching character unless the regular expression you specify must look for an event that contains only the pattern you are specifying; for example, REGEX="^test\$" will search for events containing the word "test" (without quotes) only. Special regular expression characters such as \ and ? need to be escaped.
Secondary Delimiters: = . : / \ @ - ? # \$ & - %	You can also search for keywords containing secondary delimiters once you have configured the full-text search options as described in "Full-text Search Options" on page 322 . Example: You can search for hp.com in a URL http://www.hp.com/ap ps by specifying hp.com as the search string.	You can search for these characters. Enclose value in double quotes if value contains any of these characters. Example: name="John."	<ul style="list-style-type: none"> Cannot contain ^ in the beginning and \$ at the end as a matching character unless the regular expression you specify must look for an event that contains only the pattern you are specifying; for example, REGEX="^test\$" will search for events containing the word "test" (without quotes) only. Special regular expression characters such as \ and ? need to be escaped.
Syntax	keyword1 boolean_operator keyword2 boolean_operator keyword3...	field_name operator field_value (List of fields in the "Event Field Name Mappings" on page 587 section.) (List of operators in the "Field-based Search" on page 80 section.)	REGEX="<REGEX1>" REGEX="<REGEX2>" ..

Table 4-6 Query Syntax Requirements (Continued)

Behavior	Full Text Search	Field Search	Regular Expression
Tab Newline { " *	Cannot search for these characters. Examples: "John{Doe" is invalid	No restrictions. Enclose special character in double quotes. Escape the wildcard character and double quotes. Example: name="John* \"Doe" (matches John* "Doe)	No restrictions. Special regular expression characters such as (,), [], { }, ", , and * need to be escaped.
Time format, when searching for events that occurred at a particular time	No specific format. The query needs to contain the exact timestamp string. For example, "10:34:35". Note: The string cannot contain spaces. For example, "Oct 19" is invalid.	Use this format to specify a timestamp in a query (including double quotes) : "mm/dd/yyyy hh:mm:ss" OR "yyyy/mm/dd hh:mm:ss timezone" OR "MMM dd yyyy hh:mm:ss timezone" where mm=month dd=day yyyy=year hh=hour mm=minutes ss=seconds timezone=EDT, CDT, MDT, PDT. MMM=First three letters of a month's name; for example, Jan, Feb, Mar, Sep, Oct, and so on. Use the <= and >= operators to narrow down the time range. Do not use = or !=.	No restrictions.
Wildcard	* Cannot be the leading character; only a suffix or in between a keyword. Examples: <ul style="list-style-type: none"> *log is invalid log* is valid lo*g* is valid 	* Can appear anywhere in the value. Examples: name=*log (searches for ablog, blog, and so on.) name="*log" name=*log (both search for *log)	* Can appear anywhere.

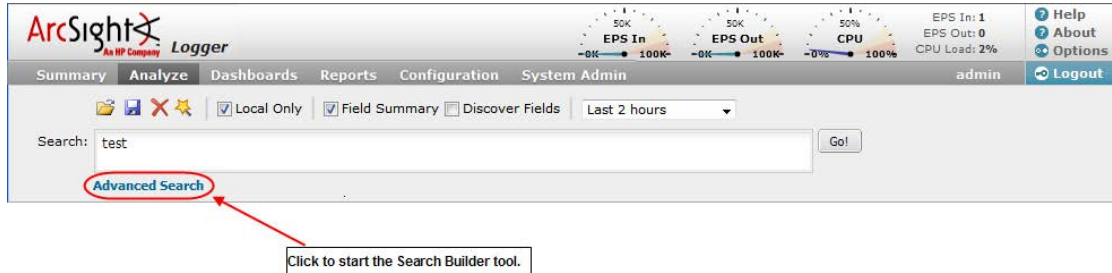
Using the Advanced Search Builder Tool

The Advanced Search tool is a Boolean-logic conditions editor that enables you to build search queries quickly and accurately. The tool provides a visual representation of the conditions you are including in a query. You can specify keywords, field-based conditions, and regular expressions using this tool. You can also specify search constraints such as peers, device groups, and storage groups (see ["Constraints" on page 92](#)). This section describes how to use the tool.

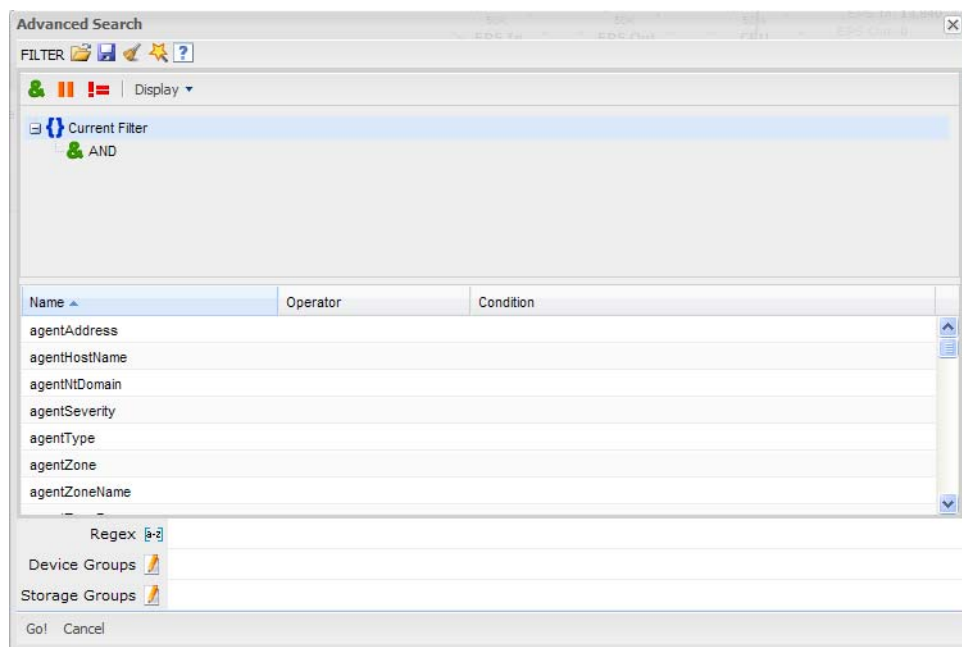
Accessing the Advanced Search Builder

To display the Advanced Search builder:

Click **Analyze > Search** to open the search page, and then click **Advanced Search**, below the Search text box, as shown in the following figure.




The Advanced Search builder is displayed, as follows:

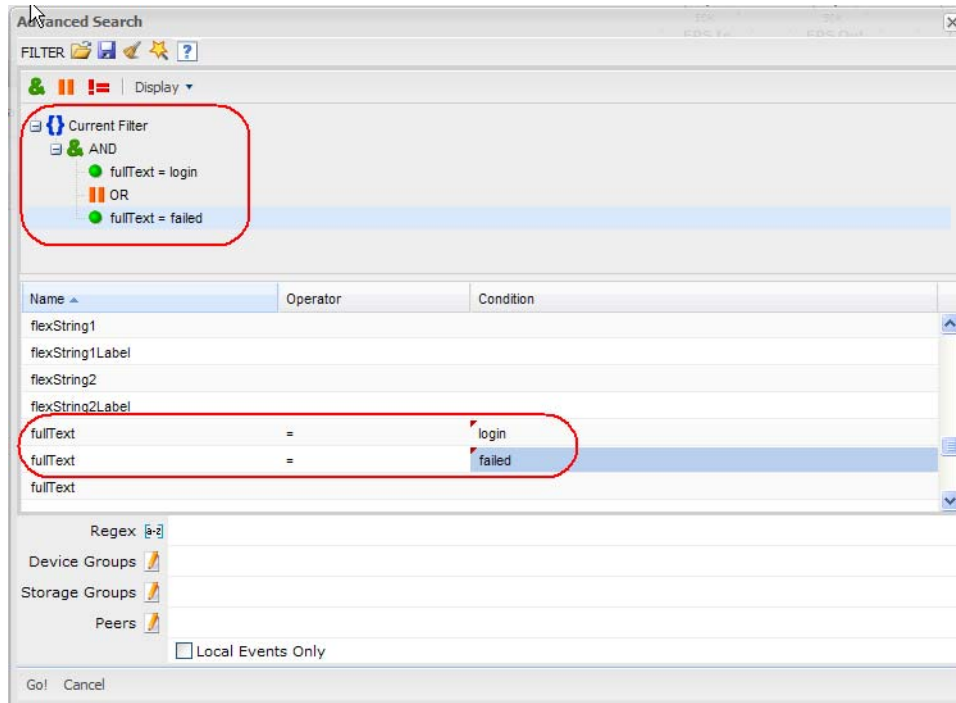


To build a new search query in the Advanced Search builder:

- 1 Click **Analyze > Search** to open the search page, and then click **Advanced Search**.
- 2 Select the Boolean operator that applies to the condition you are adding from the top of Search Builder. You can select these operators:

Operator	Meaning
&	AND
	OR
!=	NOT

- 3 If you want to load a system or saved filter, or a saved search, click the  icon. Select the filter or the saved search from the displayed list and click **Load + Close**.
- 4 For more information, see [“Saving Queries \(Saved Filters and Searches\)”](#) on page 126 and [“System Filters/Predefined Filters”](#) on page 128.
- 5 To add a keyword (full-text search) or field condition:
 - a Locate the field you want to add under the Name column.
To specify a keyword (full-text search), use the *fullText* field under the Name column, as shown in the following figure.




- b Click the Operator column associated with the field, select the operator from the displayed list, and press **Enter**.
- c Only operators applicable to a field are displayed in the list.
- d In the Condition column associated with the field, enter a value and press **Enter**.



- You cannot specify a range of IP addresses. Therefore, to search for multiple IP addresses in a range, use the CONTAINS operator and wildcard characters in the Condition column; for example, enter "192.0.2.*".
- To edit a condition, right click on the condition for a drop-down menu that enables you to edit, cut, copy, or delete the condition.

- 6 Repeat [Step 1](#) through [Step 5](#) until you have added all the conditions.
- 7 If your search query will also include a regular expression, type it in the Regex field.

- 8 If you want to constrain your search query to specific device groups, storage groups, and Loggers, click the  icon next to the constraint category. Select the relevant groups and Loggers. (To select multiple groups, hold the Ctrl-key down.)

You can specify devices or device groups in the Device Groups constraint.


The Logger constraint category is displayed only if Loggers are configured on your Logger.

If multiple values are selected for a constraint, those values are OR'ed together. For example, if you specify Device Group A, B, C, the query will find events in Device Group A, B, or C.

- 9 Click **Go**.

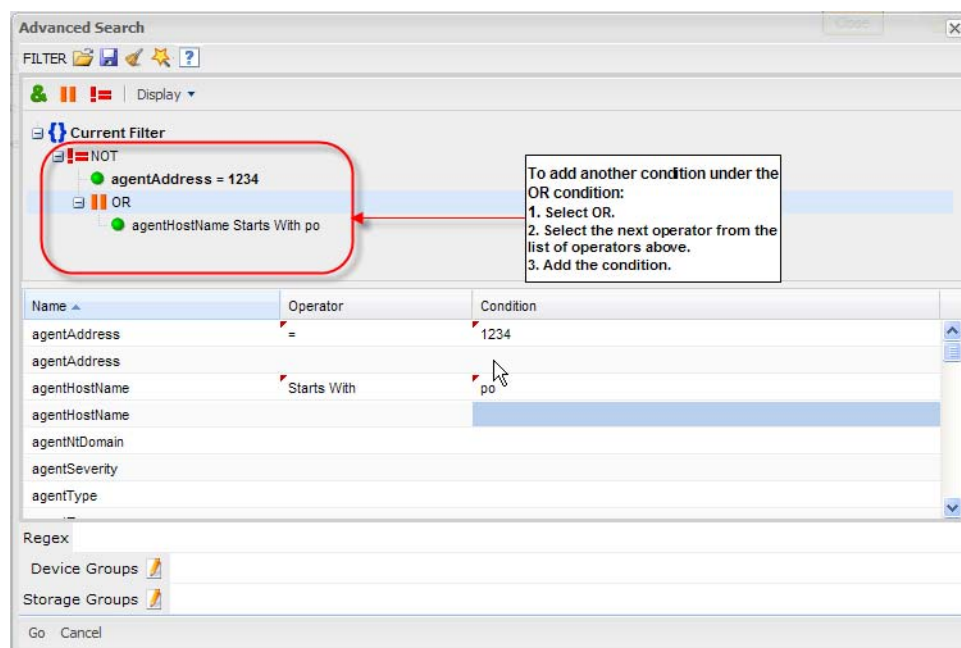
The query is automatically displayed in the Search text box and is ready to be run.

OR

Click the  icon to save the query (referred as Saved Filter or a Saved Search) for a later use. For more information about saving queries, see ["Saving Queries \(Saved Filters and Searches\)" on page 126](#).

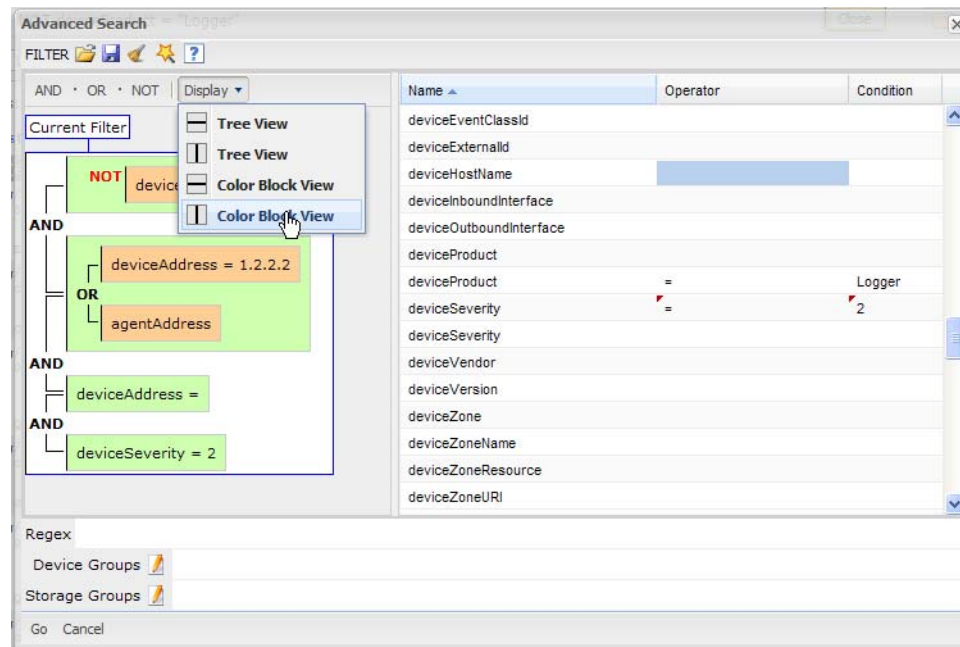
Nested Conditions

You can create search queries with nested conditions in Search Builder. To do so, click the operator under which you want to nest the next condition and add the condition as described in ["Accessing the Advanced Search Builder" on page 97](#).



Alternate Views for Query Building in Search Builder

By default, a tree view representation of the conditions is displayed, as shown in the previous figures in this section. You can change the view to a color-block scheme and also adjust whether the fields you select are displayed in the lower part of the screen or to the right of where conditions are displayed, as shown in the following figure.




To change views:

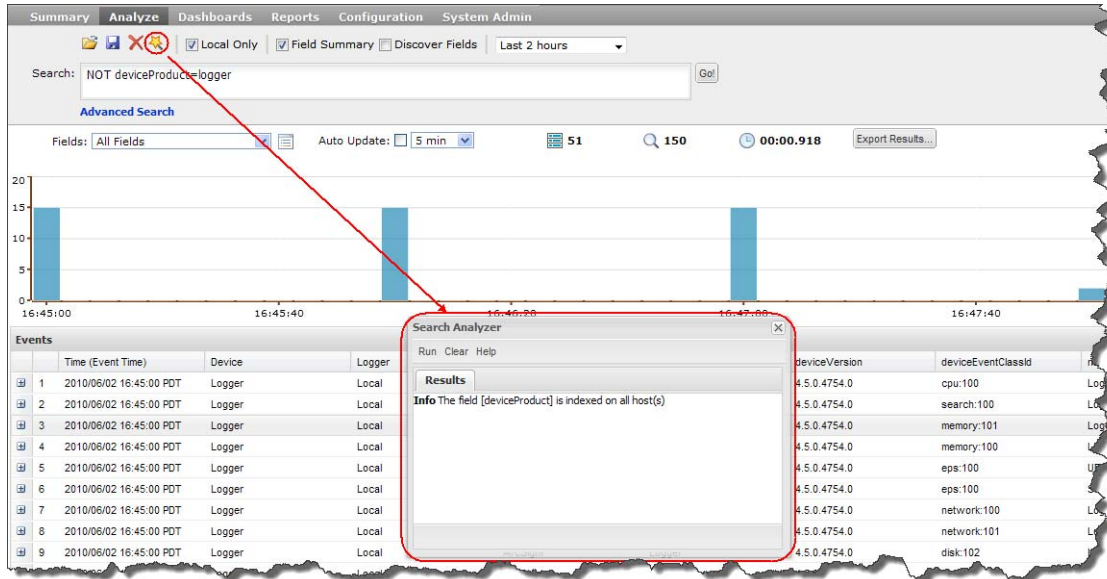
Click **Display** in the Search Builder tool and select the view of your choice.

Search Analyzer

A query's performance is dependent on many factors such as load on the system, size of data to be searched, indexed or non-indexed fields included in the query, the complexity of a query (a large number of conditions, wildcard characters, nesting), and so on.

The Search Analyzer tool analyzes a query to determine if any of the fields included in the query are non-indexed for the time range specified and thus impact the query's performance.

You can run this tool as needed; for example, if a query runs slower than expected. You can use Search Analyzer on a query after you have run it or while building a query using the Search Builder. Click the  icon to access the Search Analyzer tool, as shown in the following figure.



Performance Optimizations for Indexed Fields in Search Queries

Even though a search query includes indexed fields, you might not realize the performance gain you expect in these situations:

- When you include indexed and non-indexed fields in a query. Therefore, HP recommends that you identify the fields that you will most commonly use in queries and index all those fields.
- When you fields that are not super-indexed or field operators other than = in a needle-in-a-haystack search, your search speed may not see the expected performance increase for super-indexed fields. For fastest results when searching for rare values, be sure to follow the recommendations in [“Searching for Rare Field Values” on page 83](#).
- When you perform search on data in a time range in which a currently indexed field (included in the query) was non-indexed.
- For example, you index the “port” field on August 13th at 2:00 p.m. You run a search on August 14th at 1:00 p.m. to find events that include port 80 and occurred between August 11th and August 12th. The “port” field was not indexed between August 11th and the 12th; therefore, the query runs slower.
- When you include a field in your search query that Logger is in the process of indexing. Therefore, allow some time between adding a field to the index and using it in a search query.
- When a query that includes indexed field is performed on archived events, the query runs slower than when the data was not archived. This occurs because the index data on Logger is not archived with events.

Regex Helper Tool

The Regex Helper tool enables you to create regular expressions that can be used with the `rex` pipeline operator to extract fields of interest from an event. (For information about `rex`, see [“Search Operators” on page 86](#) or [Appendix B, Using the Rex Operator, on page](#)



553.) This tool not only simplifies the task of creating regular expressions for the `rex` operator but also makes it efficient and error free.

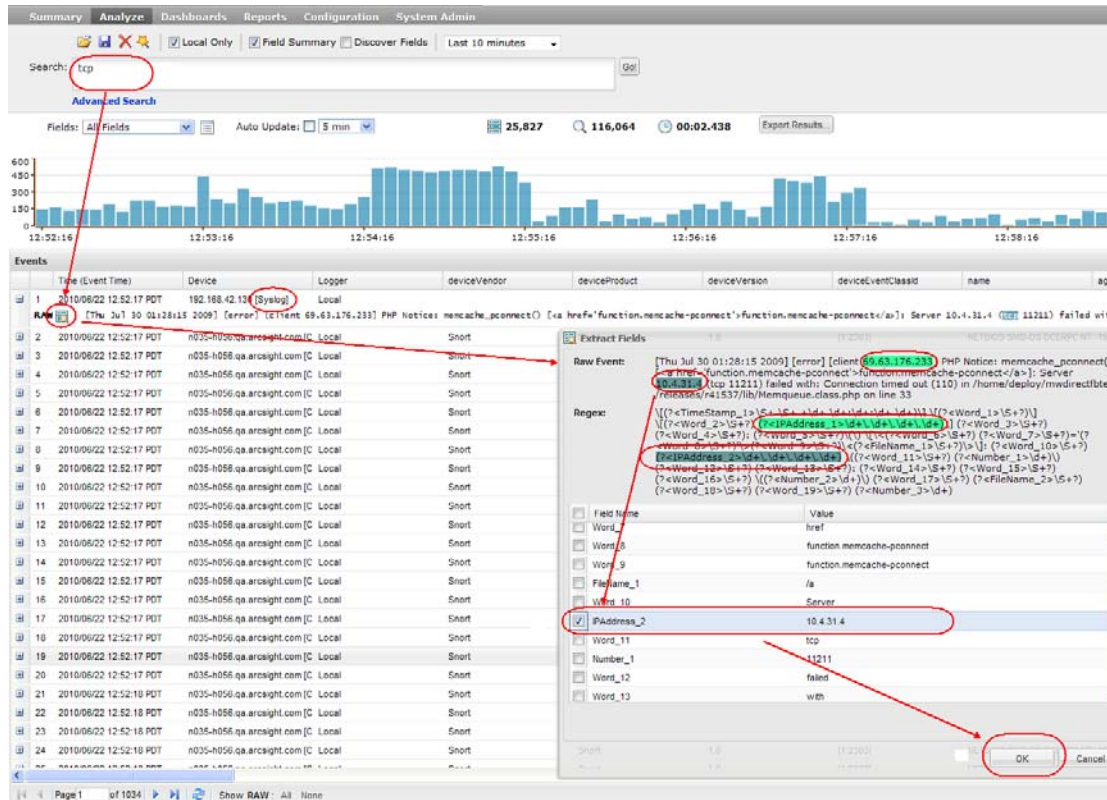
The tool, which is only available for non-CEF events (unstructured data), parses a *raw syslog event* into fields and displays them as a list. You select the fields that you want to include in the `rex` expression of a query. The selected fields are automatically inserted in a search query as a `rex` expression.

To use the tool, you need to perform the following steps:



These steps are also depicted in the figure that follows the steps.

- 1 Enter a search query that finds events of interest to you. (For information about running a search, see [“Searching for Events on Logger”](#) on page 106.)
- 2 Identify a syslog event that you want to analyze further. For example, in the shown figure, event #1 is the event we will analyze further.
- 3 Click the  icon (in the left-most column) for the identified event to expand it and display its raw event.
- 4 Click the  icon (next to the word **RAW**) to launch the Regex Helper tool.
- 5 Select the fields that you want to extract.
- 6 Click **OK**.



The screenshot shows the Logger Administrator's Guide interface. At the top, there is a navigation bar with tabs: Summary, Analyze, Dashboards, Reports, Configuration, and System Admin. Below the navigation bar, there is a search bar with the text "Search: ttp" and a "Go" button. To the right of the search bar, there are filters: "Fields: All Fields", "Auto Updates: 5 min", and a status bar showing "25,827" events, "116,064" fields, and a time range of "00:02:438".


Below the search bar, there is a bar chart showing event counts over time. Below the bar chart, there is a table of events. The table has columns: Time (Event Time), Device, Logger, deviceVendor, deviceProduct, deviceVersion, deviceEventClassid, name, and age. Event #1 is highlighted, showing a timestamp of "2010/06/22 12:52:17 PDT", device "192.168.42.13 [Syslog]", and logger "Local".

On the right side of the screen, the "Extract Fields" dialog box is open. It shows the "Raw Event" text: "[Thu Jul 30 01:28:15 2009] [error] [client 69.63.176.233] PHP Notice: memcache_pconnect() [a href='function:memcache-pconnect'>function:memcache-pconnect/a>]: Server 10.4.31.4 (11211) failed with: Connection timed out (110) in /home/depot/memcachedbter/redis/redis41537/lib/MemQueue.class.php on line 33". The "Regex" field contains a complex regular expression. Below the regex, there is a list of fields extracted from the event, including "Field Name", "Value", "Word_1", "Word_2", "Word_3", "Word_4", "Word_5", "Word_6", "Word_7", "Word_8", "Word_9", "Word_10", "Word_11", "Word_12", "Word_13", "Word_14", "Word_15", "Word_16", "Word_17", "Word_18", "Word_19", "Word_20", "Word_21", "Word_22", "Word_23", "Word_24", "Word_25", "Word_26", "Word_27", "Word_28", "Word_29", "Word_30", "Word_31", "Word_32", "Word_33", "Word_34", "Word_35", "Word_36", "Word_37", "Word_38", "Word_39", "Word_40", "Word_41", "Word_42", "Word_43", "Word_44", "Word_45", "Word_46", "Word_47", "Word_48", "Word_49", "Word_50", "Word_51", "Word_52", "Word_53", "Word_54", "Word_55", "Word_56", "Word_57", "Word_58", "Word_59", "Word_60", "Word_61", "Word_62", "Word_63", "Word_64", "Word_65", "Word_66", "Word_67", "Word_68", "Word_69", "Word_70", "Word_71", "Word_72", "Word_73", "Word_74", "Word_75", "Word_76", "Word_77", "Word_78", "Word_79", "Word_80", "Word_81", "Word_82", "Word_83", "Word_84", "Word_85", "Word_86", "Word_87", "Word_88", "Word_89", "Word_90", "Word_91", "Word_92", "Word_93", "Word_94", "Word_95", "Word_96", "Word_97", "Word_98", "Word_99", "Word_100".


The "OK" button is highlighted in the bottom right corner of the dialog box.

The rex expressions pertaining to the selected fields are automatically entered in the Search query box, as shown in the following figure. In the previous example, the client and server IP addresses need to be extracted from events. Therefore, IPAddress_1 and IPAddress_2 fields were selected in the Regex Helper tool. (The Regex Helper tool assigns incremental labels if a data type appears more than once in an event. For example, IP addresses are assigned IPAddress_1, IPAddress_2, IPAddress_3, and so on labels.)

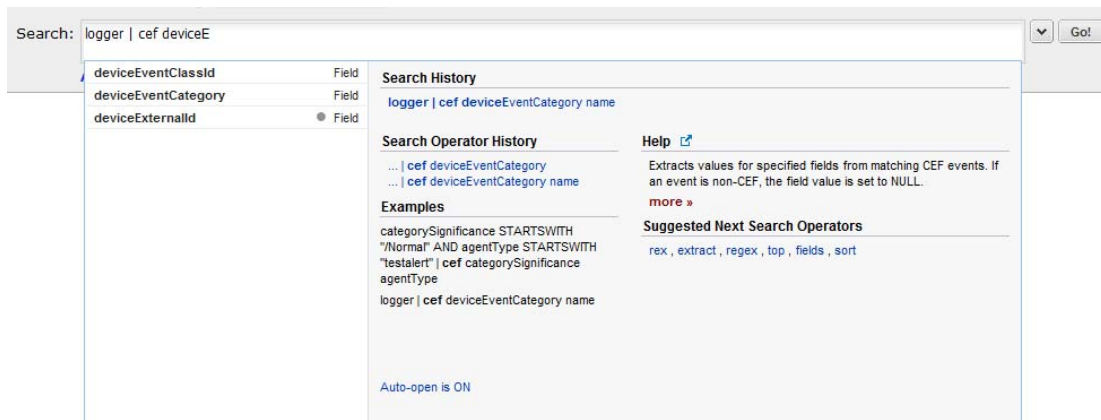
Once the two IP addresses are selected and you click **OK**, the `rex` expression that includes the regular expression for those IP addresses is displayed in the Search text box, as shown in the following example.



From this point, you can include additional pipeline operators in this query to create charts, identify the top five IP addresses, and so on. In the following example, the above query is modified to identify the top IP addresses.



Search Helper



Search Helper is a search-specific utility that automatically displays relevant information based on the query currently entered in the Search text box.

Search Helper is available by default; if you do not want the Search Helper to display information automatically, click the "Auto-open is ON" link (in the Search Helper window). The link toggles to "Auto-open is OFF". To access Search Helper on demand (once it has been turned off), click the down-arrow button to the right of the Search text box.

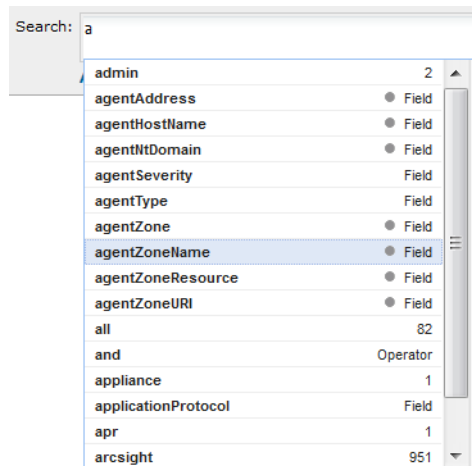
Search Helper includes following types of information:

Autocomplete search	Suggested next operators
Search history	Examples
Search operator history	Help

Autocomplete Search

The autocomplete functionality provides full-text keywords and field suggestions based on the text currently entered in the Search box. The suggestions enable you to select keywords, fields, field values, search operators, or metadata terms from a list instead of typing them in, thus enabling you to build a query expression more quickly.

When you start typing, the suggestion list displays many types of entries.



If the entered text is contained in both full-text keywords and schema fields, all of them are displayed in the suggested list.

If you type "|" (the pipeline character), the list of operators available on Logger are displayed.

The full-text keyword suggestions are obtained from the full-text keywords that are already indexed on your Logger.

If the entered Logger schema field is indexed on Logger, field values associated with it are displayed. However, if the field is not indexed, no field value suggestions are provided. The fields that are indicated by a dot (●) next to the word "Field" in the autocomplete list are **not** indexed on Logger.

The full-text keywords and field values display a count next to each suggestion that indicates the number of the instances of the keyword or field value stored on Logger.

The count represents the number of values stored for a field. The count is dependent on many factors and may not be exact. It does not indicate how many events might match the query. Many factors determine the number of event matches, including the time range, search constraints, and search operators for the query.



Note

- The autocomplete suggestions and counts are based on data stored on the local system only. Peer data is not included.
- Autocomplete suggestions and counts are reset when the Logger restarts.

Search Group filters (that restrict privileges on storage and device groups) are not enforced on the autocomplete list. Therefore, the list includes keywords, fields, field values, and counts of events in storage and device groups to which a user might not have privileges.

When an archive is loaded back on Logger, the autocomplete list does not list the full-text keywords or field values that were available when the events were not archived. This happens because index data is not archived along with the event data; therefore, when the event data is loaded back from an archive, this data is treated as unindexed.

On a Logger that is upgraded from 5.2 Patch 1 or earlier, the autocomplete list contains keywords and fields that were indexed after the upgrade; keywords and fields included in the index prior to the upgrade are not included. Therefore, if your query matches events that were in Logger prior to the upgrade, there will be an inconsistency between the displayed count and the number of events found.

To use an autocomplete suggestion:

Click the suggestion to move it up to the Search box. Then click **Go!** to run that search or continue typing in the search box to narrow your search further.

Search History

The search history displays recently run queries that match the currently entered search. Click a recent query to run it again.

The screenshot shows the Logger search interface. The search box contains the text 'logger'. Below the search box, a list of search operators is displayed, each preceded by a vertical bar (|). The operators are: | cef, | chart, | eval, | extract, | fields, | head, | keys, | rare, | regex, | rename, | replace, | rex, | sort, | tail, | top, and | where. The list is enclosed in a red rectangular box. To the right of the operators, the word 'Operator' is repeated for each entry. Below the operators, there is a section titled 'Examples' with the text: 'error alert', 'message CONTAINS "Between"', '(name="John Doe" OR name="Jane Doe") AND message="success"', and 'more »'. Below the examples, there is a section titled 'Suggested Next Search Operators' with the text: 'cef , rex , extract , regex'. Below the suggestions, there is a table with the following columns: 'Auto-open is ON', 'Logger', 'Local', 'ArcSight', 'Logger', and '5.1.0.5775.0'. The table contains three rows of data.

Auto-open is ON	Logger	Local	ArcSight	Logger	5.1.0.5775.0
4	2011/04/13 12:25:00 PDT	Logger	Local	ArcSight	5.1.0.5775.0
5	2011/04/13 12:25:00 PDT	Logger	Local	ArcSight	5.1.0.5775.0
6	2011/04/13 12:25:00 PDT	Logger	Local	ArcSight	5.1.0.5775.0

Search Operator History

Displays the fields used previously with the search operator that is currently typed in the Search text box. The Search Operator History only displays if you have previously used the operator you have currently typed to perform searches on this system. Click the operator to add it to your search.

Examples

Lists examples relevant to the latest query operator you have typed in the Search text box.


Usage

Provides the syntax for the search operator.

Suggested Next Operators

List of operators that generally follow the currently typed query. For example, if you type `logger |`, the operators that often follow are `rex`, `extract`, or `regex`. You can select one of the listed operators to automatically append to the currently typed query in the Search text box. This list saves you from guessing the next possible operators and manually typing them in.

Help

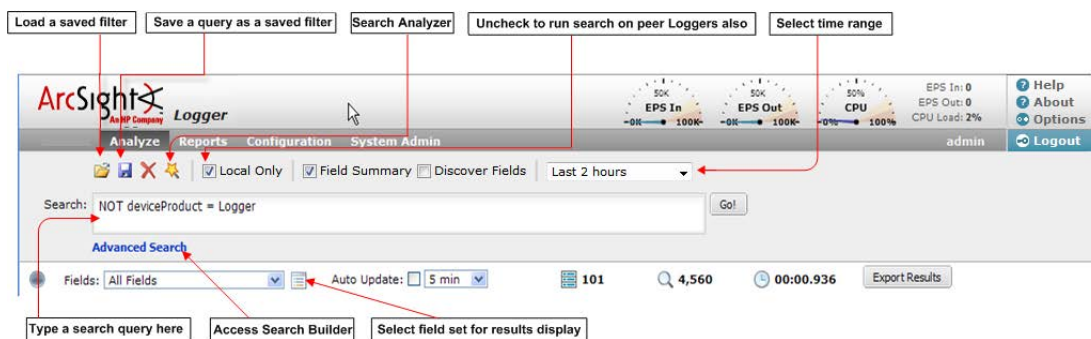
Provides context-sensitive help for the last-listed operator in the query that is currently typed in the Search text box. Additionally, if you click the  icon, Logger online Help is launched.

Searching for Events on Logger

To perform local searches, a user must belong to a Logger Search Group with the “Search for events” user right set to Yes.

To perform searches on peers and view the search results, a user needs to belong to these user groups with the listed permissions:

- ◆ Logger Search Group with “Search for events on remote peers” user right set (checked).
- ◆ Logger Rights Group with the “View registered peers” user rights set (checked).



To search for events on Logger:

- 1 Click **Analyze > Search**.
- 2 Use the following default values or change them suit your needs:
 - a **Local Only:** When peers have been configured for your system, the Local Only checkbox will display. Local Only is checked by default. If you want to include peers in your search, uncheck the Local Only checkbox. If you do not see this checkbox, no peers have been configured.
 - b **Time Range:** By default, the query is run on the data received in the last two hours. Click the drop-down list to select another predefined time range or specify a custom time range. For more information about time ranges, see [“Time Range” on page 86](#).
 - c **Fieldset:** By default, all fields (All Fields) are displayed in the search results. However, you can select another predefined fieldset or specify a customized fieldset. For more information about fieldsets, see [“Fieldsets” on page 88](#).

- 3 Specify a query expression in the Search text box using one or more of the following methods.




Refer to [“Query Expressions” on page 78](#) for a list of exceptions and invalid characters before you create a query expression.

- a Type the query expression in the Search text box. For information about building a query expression, including lists of applicable operators, see [“Elements of a Search Query” on page 78](#).
- b When you type a query, Logger’s Search Helper enables you to quickly build a query expression by automatically providing suggestions, possible matches, and applicable operators. See [“Search Helper” on page 103](#) for more information.
- c Use these guidelines to include various elements in a search query:
 - For a complete list of fields in Logger schema, see [“Indexing” on page 121](#).
 - Metadata terms (`_storageGroup`, `_deviceGroup`, `_peerLogger`)
Type “_s” (for storage group), “_d” (for device group), or “_p” (for Logger) in the Search text box to obtain a drop-down list of constraint terms and operators.
 - Regular expression term (`|REGEX=`)




If your query expression includes multiple device groups and storage groups to which search should be constrained, make sure that the group names are enclosed in a square bracket; for example, `_storageGroups IN [“SGA”, “SGB”]`.

- Click **Advanced** to use the Search Builder tool. (See [“Using the Advanced Search Builder Tool” on page 96](#) for more information.) Also, use this option to specify device groups, storage groups, and Loggers to which search should be limited.
- d Click the  icon to load a saved filter, a system filter, or a saved search. Select the filter or the saved search from the displayed list and click **Load+Close**.

For more information, see [“Saving Queries \(Saved Filters and Searches\)” on page 126](#) and [“System Filters/Predefined Filters” on page 128](#).
- 4 Click **Go**.

The search results are displayed in the bottom section of same screen in which you ran the search. For more information about how search results are displayed and various controls available within the user interface to use those results, see [“Understanding the Search Results Display” on page 109](#).

You can also save the search you ran as a saved filter or saved search. Click the  icon to do so. For more information about a saved filter or a saved search, see [“Saving Queries \(Saved Filters and Searches\)” on page 126](#).

Advanced Search Options

The advanced search options enable you to tune search operations to suit your environment. The options are discussed in [“Tuning Advanced Search Options” on page 321](#).

Searching Peer Loggers (Distributed Search)

When you run a search query, by default, only your local Logger is searched for matching events. However, when specifying a query, you can select an option to run the search on the peer Loggers. You can also select the Loggers to which the search should be constrained, as described in [“Searching for Events on Logger” on page 106](#).

Follow these guidelines for searching across peers:

- Searches across peers are limited by the ability of the earliest version peer. For example, search speed is limited by the speed of the earliest peer.
- If Loggers do not have identical storage or device group names, a search query operation skips searching for events for those groups on those peers.
- If you added custom schema fields to your Logger schema, those fields must exist on all peers. Otherwise, a search query containing those fields will not run (when run across peers) and return an error. See [“Adding or Importing Schema Fields” on page 345](#).
- A user needs to belong to these user groups with the listed permissions set to perform searches and view their search results:
 - ◆ Logger Search Group with “Search for events on remote peers” user right set (checked).
 - ◆ Logger Rights Group with the “View registered peers” user rights set (checked).
- When a Logger becomes unavailable during a search operation, error messages are displayed. The displayed message varies depending on the error detected. For example, you may see the following message:

```
[Logger IP address] Error: Get Query Statistics  
[Logger IP address] Error: Remote exception (does not authorize  
the request. Please check if remote has relationship with your  
logger)
```

This error message occurs when the Logger cannot be reached. To fix the issue, restore the relationship and run the search again. The error messages may still display for the search that was in progress even after the relationship is restored. However, you can ignore such messages if they go away when you run a new distributed search.

For more information about peers, see [“Peer Loggers” on page 326](#).

Tuning Search Performance

Search performance depends on many factors and will vary from query to query. Some of factors that can impact search performance are listed below. To optimize search performance, ensure that you follow these recommendations:

- When searching for uncommon field values, use super indexing to narrow the range of data that needs to be searched, as described in [“Searching for Rare Field Values” on page 83](#).
- Enable field-based indexing for all fields that occur in your events. When events are indexed, Logger can quickly and efficiently search for relevant data. By default, a recommended set of fields are indexed on your Logger; you might need to add additional fields, as described in [“Indexing” on page 121](#).
- Avoid specifying a time range that results in a query that needs to scan multi-millions of events.
- Limit the search to specific storage groups and peers.

- Reduce other load on the system when your query needs to run, such as scheduled jobs, large number of incoming events, and multiple reports being run.
- Before running a query, make sure all Loggers on which it will run support the query features.

For more information on improving search performance, refer to the *Logger Configuration and Tuning: Best Practices* technical note.


Understanding the Search Results Display

After you have initiated a search, the search results are displayed in the bottom section of the same screen in which you ran the search.





While the search is in progress, the Go! button changes to Cancel—click Cancel to terminate the search early. When a query is running, search results are displayed as matching events are found. Therefore, when you click Cancel, any matching events found so far are displayed as the search results. This facility might be helpful in cases when the query needs to scan a large data set, but the search results displayed so far display the events you were looking for. You can further process the displayed (partial) results; for example, export the results, use the histogram to drill-down the results, or click on any text in the Search Results to add it to the query for further drill-down of the search results.



If a query includes chartable operators such as chart, rare, or top, and the query is terminated early, a chart of the partial results is not displayed. Additionally, if a query includes the head, tail, or sort operators, partial results are not generated.

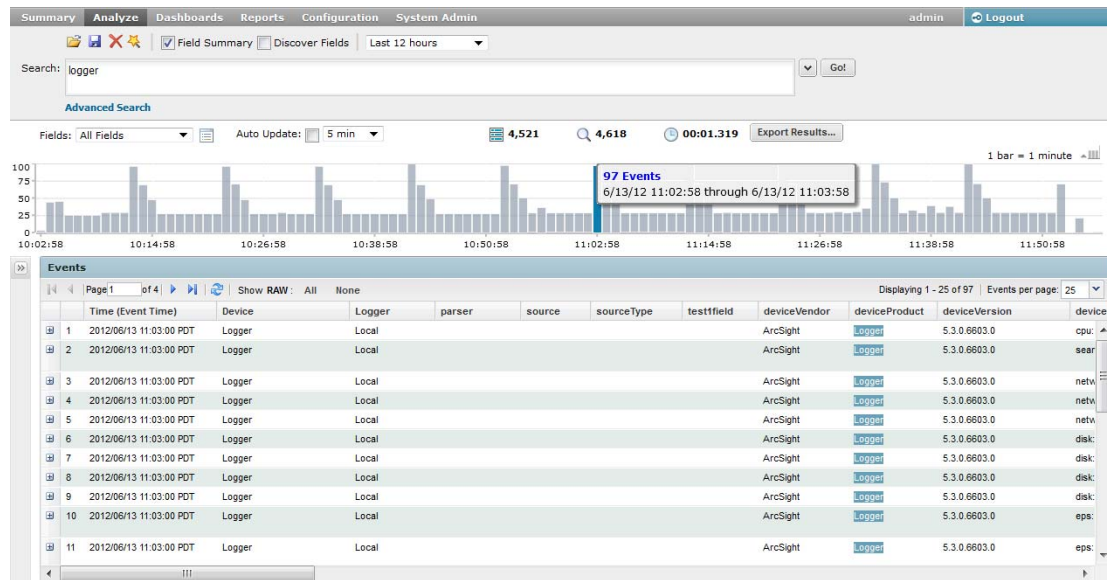
A search operation can take time when millions of events need to be searched. When the first screen of events that match the specified conditions is available, Logger automatically pauses the search and displays the matched events. By default, 25 events are displayed on one screen. Event data is categorized by field name with each field displayed as a separate column, as shown in the following figure. For example, time when the event was received on the Logger (Event Time) is displayed under Time (Event Time). Each event is also available in its raw form and can be viewed by clicking the  icon in the left most column.

To see all raw events, click **All** at the top of the Search Results display. To collapse raw events, click **None**. The column width for each column is adjustable.

To see the next screen of events, click ; or  to go to the last page. Once you are past the first screen of events, you can click  to go back to the previous screen; or  to go to the first page.

To change the number of events displayed per screen, open the Events per Page drop down menu and select the number of events to display.

The Search Results page displays a histogram that provides a graphical representation of the events that match a search query. The distribution is based on the time range specified in the query. That is, the X-axis represents event time and Y-axis represents the number of matching events, as shown in the following figure.



You can drill-down to events in a specific time period by clicking the histogram bar representing the time period. If you mouse over a bar in the histogram, the number of events scanned and number of events matching the query and the time it took to run the search is displayed.

Below the histogram, events are shown in table form, one row per event. Terms that match your query are highlighted in blue to make it easy to see why an event matched the query.

To view the raw event of a listed event, click the icon to the left of the matching event. You can also view the Syslog raw events in a formatted column called rawEvent if you have enabled the "Populate rawEvent field for syslog events" option on the Search Options page, as discussed in ["Tuning Advanced Search Options" on page 321](#). Also, see ["Fieldsets" on page 88](#) to learn more about the rawEvent field.

As you roll the mouse over other terms in the events table, they highlight in green. The user interface allows you to drill-down into the displayed search results by clicking a green-highlighted term to add it to the current query. For example, if you search for "login" and roll over the word "fail" in the search results, "fail" will highlight in green. Click the word "fail" to change the query to "login AND fail." You can also highlight and copy text from any displayed column. This feature is handy when you need to copy an IP address or a URL. (Highlight the term by scrolling over it. Then, right-click your mouse to display the Copy option.) You can select any fields from the search results. Search results are sorted by receipt time.

Use these keyboard shortcuts to select terms from the displayed search result columns or the raw events to refine your search query:

- Click the term in search results to add the selected term to the search query, and rerun the search.
- Ctrl+click to replace the entire search query with <field name> + "CONTAINS" + <selected term>, and rerun the search.
- Alt or Shift + click the term in search results to add NOT to the term, and rerun the query, thus eliminating the events that match the term you selected.
- You can add multiple NOT conditions by holding the Alt key and selecting terms in search results. When multiple conditions are added, they are joined by AND operators.

- You can combine Ctrl+Alt, (or Ctrl+Shift) to replace the search query with NOT + <field name> + "CONTAINS" + <selected term>.

A Field Summary panel is displayed on the left side of the matched events. This section lists the fields that occur in matching events and the number of unique values for each in those events. For more information about Field Summary, see [“Understanding Field Summary” on page 116](#).

User-defined Fields in Search Results

When a search query matches events that were received from a defined source type and were parsed using a pre-defined or user-defined parser, the search results include a parser field, and may include fields for the source type, and source, depending on the setting in the Search Options tab. For more information, see [“Tuning Advanced Search Options” on page 321](#).

The following table describes the purpose of these fields.

Field	Description
parser	Indicates whether an event was parsed or not, and which parser was used. If the event was parsed, this field contains the name of the parser. If the event was not parsed successfully, this field contains “Not parsed”. If no parser is defined for the source type or if there is no source type, the field is blank.
source type	The type of file from which the event was received, as defined on the Source Type page (Configuration (or Configuration > Settings) > Event Input > Source Types). If no source type was applied when the event was received, this field is blank. You can control whether this field is displayed from the Search Options tab.
source	The name of the log file from which the event was received. For example, /opt/mnt/testsoft/web_server.out.log. If no source was applied when the event was received, this field is blank. You can control whether this field is displayed from the Search Options tab.

User-defined fields are also created when a search query includes operators such as `rex`, `extract`, and `rename`. See [Appendix A, Search Operators, on page 529](#) for information on these operators.

These fields are displayed as additional columns in the All Fields view (of the System Fieldsets). To view only these columns, select **User Defined Fieldsets** from the System Fieldsets list.

Viewing Search Results Using Fieldsets



By default, the Search Results are displayed using the All Fields fieldset, which displays all fields contained in an event. Once you select another fieldset, it becomes your default view until you change it the next time. For a detailed discussion about fieldsets, see [“Fieldsets” on page 88](#).

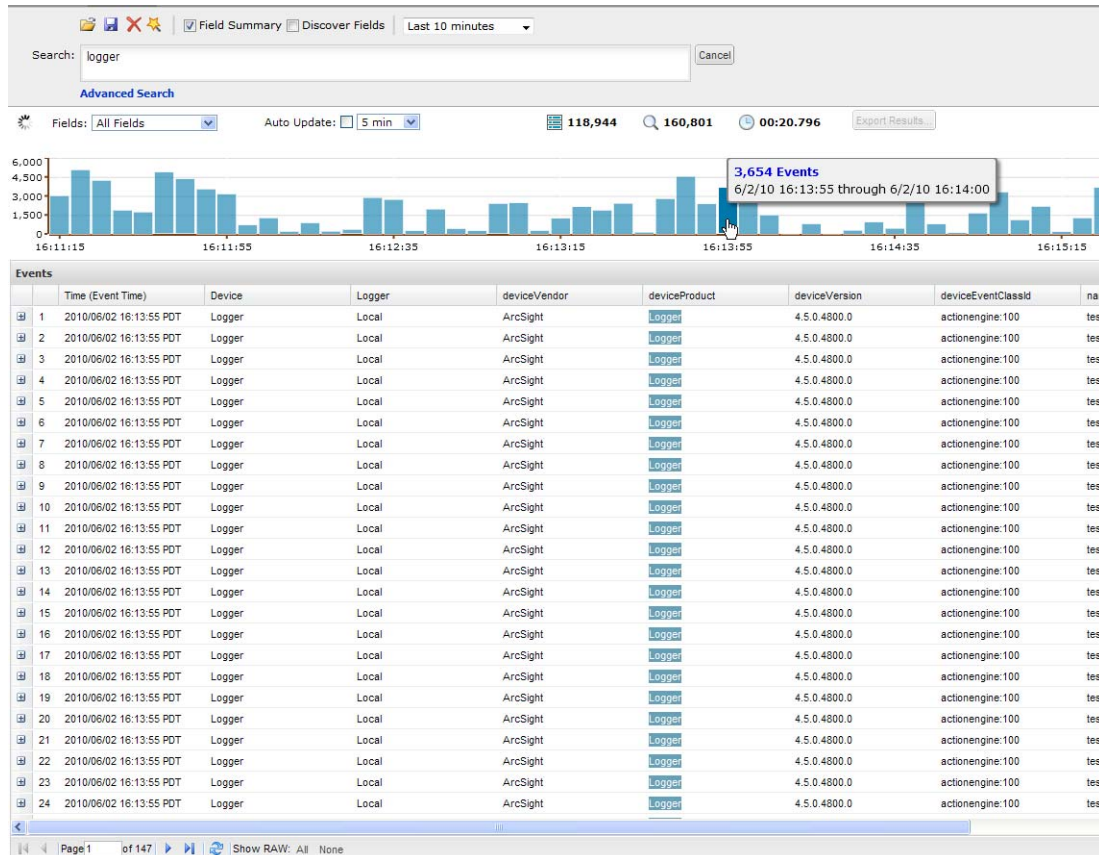
If you view the Search Results using the Raw Event fieldset, remember these guidelines:

- Even though the rawEvent column displays the raw event, this column is not added to the Logger database and is not indexed. Therefore, you can only run a keyword (full-text) or regular expression to search on the event.
- You can use the Regex Helper tool to identify strings from the raw syslog events in the rawEvent column that you want to add to a query. (You cannot use the Regex Helper for CEF events displayed in the rawEvent column.) See [“Regex Helper Tool” on page 101](#) for details about the Regex Helper tool.

Using the Histogram

Use the following guidelines to effectively and efficiently use histograms:

- Histogram of the matching events is generated automatically. You cannot disable it, however, you can click  to the upper-right corner of the histogram to hide it. To display a hidden histogram, click the  icon.
- Histogram is based on the Logger receipt time of the events (similar to search queries that also use the Logger receipt time to search for events).
- The time distribution on the X-axis is determined automatically.
- You can mouse-over any histogram bar to view the number of matching events and the date and time period that the bar represents.
- You can drill-down to events in a specific time period by clicking the bar on the histogram that represents that time period. The selected section is highlighted and the events matching that time period are listed below the histogram. The histogram continues to display the distribution of all of the matching events, as shown in the following figure. For example, if you select a bar that represents 11,004 events on 2/22/2010 from 12:25:49 a.m. to 12:26:49 a.m. in the following histogram, the details of those events are listed below the histogram; however, the histogram displays all time units and the associated bars. You can also select multiple consecutive bars on the histogram to view matching events in all of the selected time units.
- To deselect a selected bar, click it.



- A histogram is progressively built and displayed as events match a search query. If the search query needs to scan a large amount of data or a large time period, the histogram displayed initially might refresh multiple times while the query is running. To view the complete (and final) histogram of a search query, wait until the query has finished running (that is, the screen does not display the circular “waiting” icon anymore).
- The time range on the X-axis might not match the time range specified in the search query because the start and end times on the X-axis are determined by the event times of the first and last matching events of the search query.
- The first one million matching events are plotted on the histogram. If a search query matches more than one million events, an informational message is displayed on the screen.
- If you need to use the histogram view for event analysis for a search query that matches more than one million events, ArcSight suggests that you adjust the time range specified in your search query such that less than one million are matched to obtain a complete and meaningful histogram or use a pipeline operator such as top, head, or chart to further refine search results such that the total number of hits is under one million events.

Multi-line Data Display

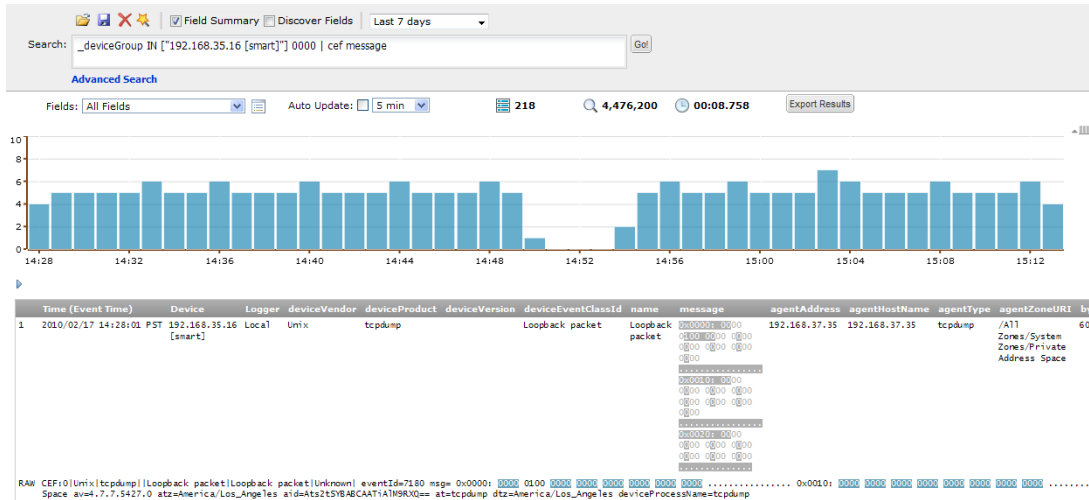
An event message might span multiple lines separated by characters such as newline (\n) or carriage return (\r). For example,

```

0x0000: 0000 0100 0000 0000 0000 0000 0000 0000 .....
0x0010: 0000 0000 0000 0000 0000 0000 0000 0000 .....
0x0020: 0000 0000 0000 0000 0000 0000 0000 .....

```

The Logger user interface displays such a message in the expected multi-line format and does not remove the line separators and collapse the message into one line, as shown in the following figure.



Auto Updating Search Results

The Auto Update feature executes the search over specified intervals, updating the search results if new events match the query.

Depending on your needs, you can auto update the search results every:

- 30 seconds
- 60 seconds
- 2 minutes
- 5 minutes (default)
- 15 minutes

You can enable this option for a search operation before or after running it. Once you enable this option, the setting persists for all search operations until you explicitly disable it.

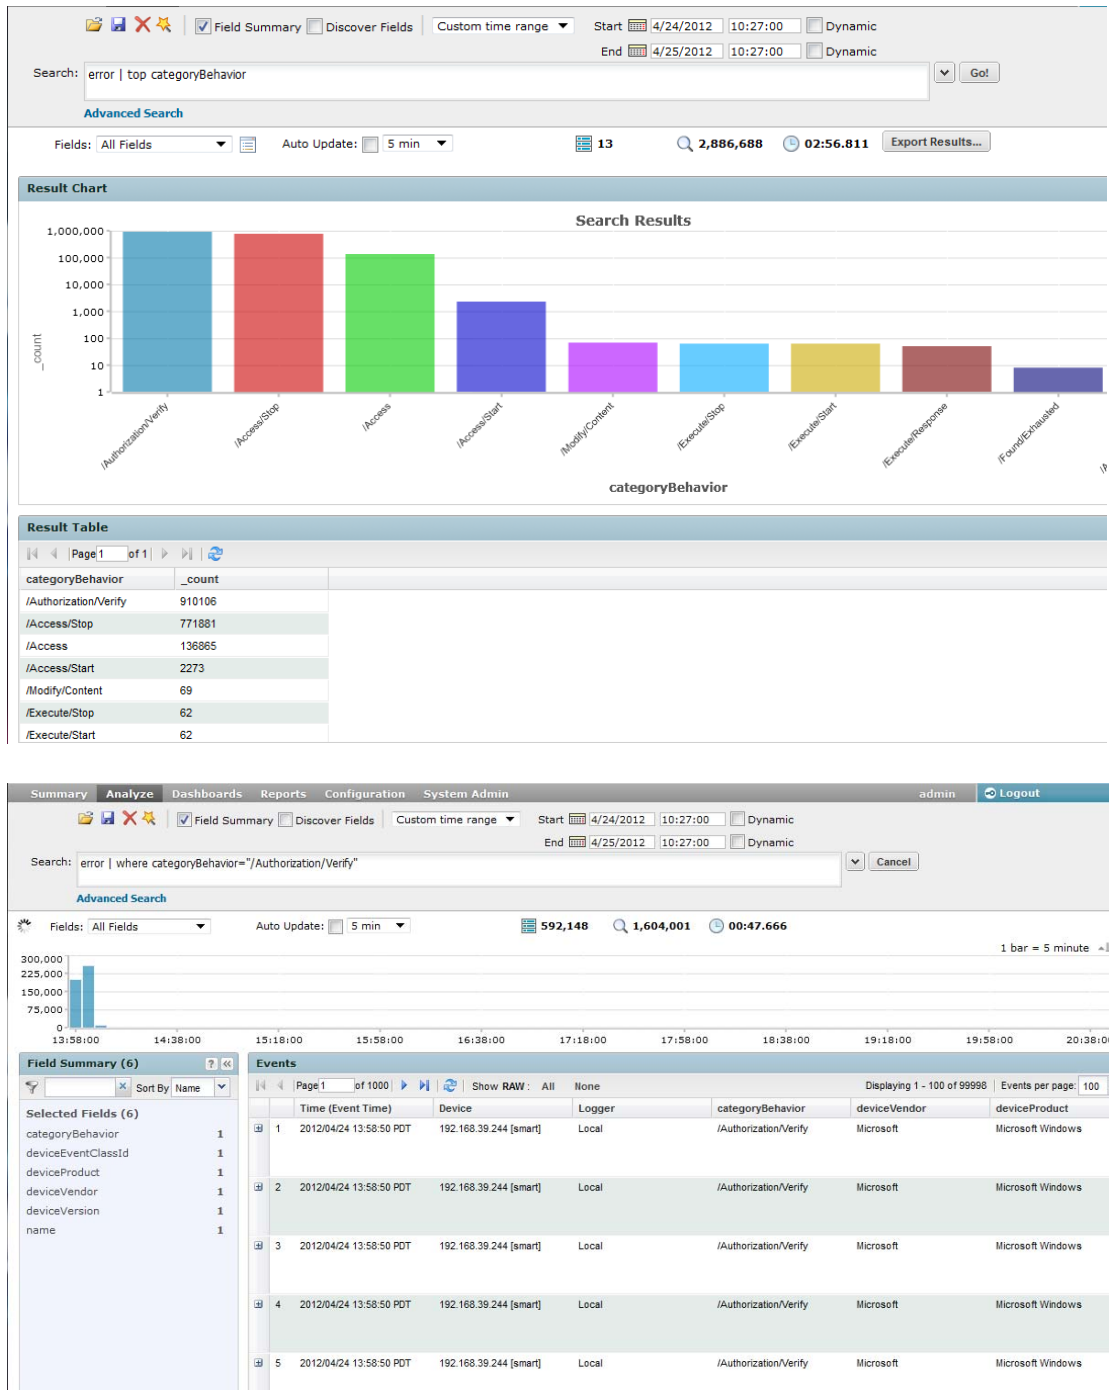
To auto update search results:

- 1 Click **Analyze > Search**.
- 2 Check the **Auto Update** box and select the refresh interval if different from the default, 5 minutes.

Chart Drill Down

Aggregated search operators such as chart, top, and rare generate charts of search results. The chart drill down feature enables you to quickly filter down to events with specific field values. You identify the value on a search results chart and click it to drill-down to events that match the value. For example, in the following chart, if you want to see events in

which the categoryBehavior field is /Access/Stop, click the second column to display events shown in the second figure.

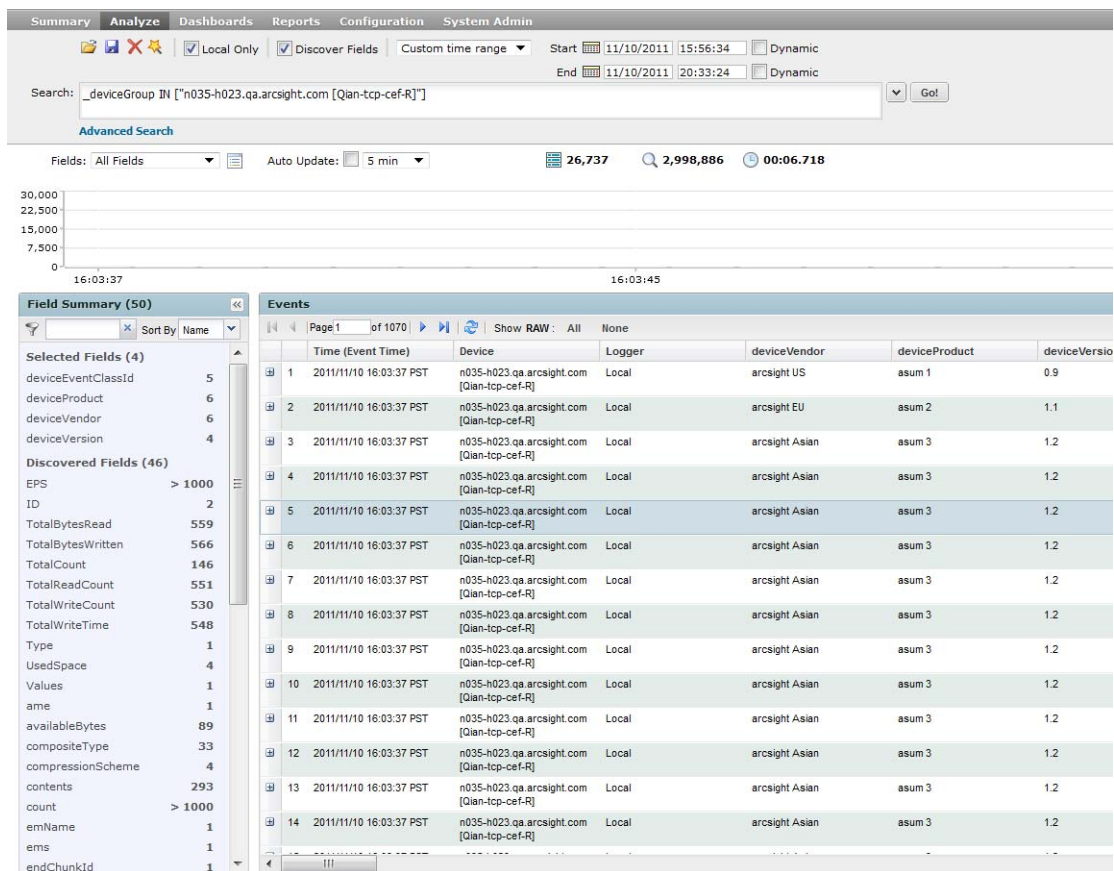


When you click on a chart value (a column, bar, or pie section), the existing search query is modified to include the WHERE operator with the field name and value, and automatically rerun. If you need to return to the original query from the drill-down screen, use the Back function of your browser.

Understanding Field Summary

When a query is run, the Field Summary panel lists the CEF and non-CEF fields that occur in matching events and the number of unique values for each in those events. This panel is only displayed for queries that do not generate charts. If a peer search is performed, the summarized field values include counts from peer Loggers.

The Field Summary panel contains two sections: Selected Fields and Discovered Fields. The Selected Fields section lists the CEF fields, while the Discovered Fields section lists the non-CEF fields discovered in raw events (described later in this section). By default, the Selected Fields list contains these fields: deviceEventClassId, deviceProduct, deviceVendor, deviceVersion, and name; you can edit this list to suit your needs, as described in [“To change the default Selected Fields list:” on page 118](#). For both lists, by default, the top 10 values for each field are listed.



The Field Summary feature can automatically discover non-CEF fields from a raw event if the Discover Fields checkbox (next to the Local Only checkbox on top of the “Search:” text box) is checked.

By default, the Field Summary feature is enabled however the Discover Fields option is disabled. If you need to enable the Discover Fields option for all searches on your Logger, change the default values (“No”) on the Search Options page (**Configuration** (or **Configuration > Settings**) > **Search Optimization** > **Search Options**) to “Yes” for these options, as shown in the following figure.

Field Summary Options

Use Field Summary: Yes

Discover fields: No

However, if you need to use the Discover Fields option occasionally—not for all searches—you can enable this option for one-time use on the user interface page from where you run the search query (**Analyze > Search**). To do so, click the Discover Fields checkbox above the Search textbox before clicking Go! to run the query. Selecting these options on the Search page overrides the setting for these options on the Search Options page.

To auto discover fields, the raw event must contain data in the “key=value” format, and none of these characters can be the first character of the “value”: comma, space, tab and semicolon. For each “key=value” pair found in a raw event, a new field of the name “key” is created. The Field Summary includes a summary of the values for all the new fields under the Discovered Fields section. The discovered fields are assigned the type “String” by default. The auto-discovery capability works only if at least 2,500 of the first 10,000 matching events contain “key=value” pairs. If this threshold is not met, auto discovery is automatically turned off. However, this threshold does not apply if there are less than 10,000 matching events; in that case, fields are discovered regardless.

You can drill-down on any of the listed fields or a specific value of the listed fields. For example, you might want to view all events containing deviceEventClassId (specific field) or you might want to view events of deviceEventClassId “storagegroup:100” (specific value of a field).

For fields whose values are of type String, you can view all events, view the top 10, or create charts of the matching events. For fields whose values are of type Numeric, you can perform mathematical operations such as average, min, and max.

Every time you run a query or drill-down on a specific field or value, a new query using the newly selected criteria is run and the Field Summary list is updated.

You can search for a specific field or filter the listed fields by specifying a filter criteria in the Filter text box located at the top of the Field Summary panel, as shown in the following figure.

Field Summary (5)

Filter: [] Sort By: Name

Selected Fields (5)	
deviceEventClassId	50
deviceProduct	1
deviceVendor	1
deviceVersion	3 ▶
name	36


Discovered Fields (0)

For example, if you want to see fields that begin with “device”, enter “device” in the Filter text box. To go back to the default list, click the icon. You can sort the field list by Name or Count. To do so, select the sort criteria from the “Sort By:” drop-down menu.

To change the default Selected Fields list:

- 1 Define or update an existing custom fieldset to include fields you want the Selected Fields list to contain. See [“Fieldsets” on page 88](#) for information on creating custom fieldsets.
- 2 Select the custom fieldset you defined to view search results.
- 3 After running a search query, if you select a different fieldset, the Field Summary panel displays the following message:

The Field Summary is out of sync with the Events table.



This message indicates that the fields listed in the Field Summary panel do not match the ones specified in the newly selected fieldset. To display the fields specified in the new fieldset, click **Update now**.

Refining and Charting a Search from Field Summary

When you click a field in the Field Summary, a dialog box labeled *<fieldname><number of values>* displays information about the field. From here, you can drill down to see more details and create a chart of the search results.

To view field details from field summary:

- 1 Click **Analyze > Search** to open the search page.
- 2 Check the Field Summary checkbox and then run a search.
- 3 Click the field name in the Field Summary.
- 4 The *<fieldname><number of values>* dialog box displays the top ten field values.
- 5 Optionally, click a field value to append it to the query and rerun the search.
- 6 To create a chart of the search results, click one of the Chart on values, such as **Values by time** or **Top values**.
- 7 The results display in a Result Chart and a Result Table.
- 8 In the Result Chart, click **Chart Settings** to adjust the chart.
- 9 Enter a useful **Chart Title**.
 - ◆ Select the **Chart Type** best suited to your data.
 - ◆ Set the **Display Limit**. The highest valid value is 100.
- 10 In the Result Table, you can use navigation buttons to move forward and backward through list of results, and refresh the search.

To create a PDF or CSV file containing the search results, click **Export Results**. For more information, see [“Exporting Search Results” on page 119](#).


Exporting Search Results

You can export search results in these formats:

- PDF—Useful in generating a quick report of the search results. The report includes a table of search results and any charts generated for the results. Both raw (unstructured data) and CEF (structured data) events, can be included in the exported report.
- Comma-separated values (CSV) file—Useful for further analysis with other software applications. The report includes a table of search results. Charts cannot be included in this format.

Data for the following time fields is exported in human-readable format: deviceReceiptTime, startTime, endTime, agentReceiptTime. For example, 2014/03/21 20:22:09 PDT.

The following is an example of a quick report generated in PDF format. The chart is displayed first, followed by a table of matched events (not shown in this example). All generated charts (including stacked charts) can be exported.



To export search results:

- 1 Run a search query.
- 2 Click **Export Results** in the top right-hand side of the search results screen.

3 Select from the following export options.

Option	Description
Save to local disk	The file is saved to a local system from which you are accessing Logger or is it sent to the browser for viewing or saving.
Export to remote location	On a Logger appliance, the file is written to an NFS mount, a CIFS mount, or a SAN system. On the software version of Logger, data is always stored in the <code><install_dir>/data/logger</code> directory. This directory can reside locally on the system running the Logger software, or on a remote storage system such as NFS or CIFS.
Save to Logger	The file is written to the Logger's local storage.
File Format	CSV , for comma-separated values file. PDF , for a report-style file that contains search results as charts and in tables. Charts are only included in the PDF file if the search query contains an operator that creates charts, such as chart, top, and so on.

Export file name	(Available only when the “Export to remote location” option is selected) Specify the name of the file to which events will be exported. If a file of the specified name does not exist, it is created. If a file of the specified name exists and the Overwrite box is not checked, an error is generated. If the Overwrite box is checked, the existing file is overwritten.
Title	(Optional, available only when the File Format is “PDF”) A meaningful name that appears on top of the PDF file. If no title is specified, “Untitled” is included.
Fields	A list of event fields that will be included in the exported file. By default, all fields are included. You can enter fields or edit the displayed fields by deselecting All Fields. To export fields created as a result of <code>rex</code> , <code>extract</code> , <code>rename</code> , or <code>eval</code> operators, or field created when a parser is applied to an event, ensure that <code>*user</code> is selected in the Fields list.

Option	Description
Chart Type (for PDF only)	(Available only when a chart is available in search results) Type of chart to include in the PDF file. You can select from: Column, Bar, Pie, Area, Line, Stacked Column, Stacked Bar. Note: If the Chart Type is different from the chart displayed on the Search Results screen, the value selected for this option overrides the one shown in the screen. Therefore, the exported PDF contains the chart you specify for this option and not the one shown on the screen.
Chart Result Limit (for PDF only)	(Available only when a chart is available in search results) Number of unique values to plot. Default: 10 If the configured Chart Result Limit is less than the number of unique values for a query, the top values equal to the Chart Result Limit are plotted. That is, if the Chart Result Limit is 5 and 7 unique values are found, the top 5 values will be plotted.
Include Summary	Include an event count in the exported search results.
Include Only CEF Events	Only include CEF events in the exported search results.
Include Base Events	Include base events in the exported search results.

4 Click **Export**.

Scheduling an Export Operation

The time it takes to export search results is proportional to the number of events being exported. Therefore, for a large number of events, HP recommends that you schedule the export operation to be performed at a later time by saving the query and time parameters as a Saved Search, and then scheduling a Saved Search Job. For more information about Saved Search jobs, see [“Scheduled Saved Search” on page 316](#).

Indexing

Logger's storage technology enables automatic indexing of events in these ways:

- Full-text indexing—Each event is tokenized and indexed. See [“Full-text Indexing \(Keyword Indexing\)” on page 122](#).
- Field-based indexing—Event fields are indexed based on a predetermined schema. See [“Field-based Indexing” on page 122](#).
- Super indexing—Certain event fields are super-indexed for finding rare field values quickly. See [“Super Indexing” on page 125](#).

How Indexing Works

Once you have initialized Logger, it starts scanning events automatically and indexing them based on these methods—full-text (keyword) and field-based. All events received after initialization are indexed for full-text search and a default set of fields is indexed for field-based search.

All events are timestamped with the receipt time when received on the Logger. The default fields are automatically indexed. For the remaining fields, Logger uses the receipt time of

an event and the time when a field was added to the index to determine whether that event will be indexed. If the receipt time of the event is equal to or later than the time when the field was added to the index, the event is indexed; otherwise, it is not.

**Note**

Indexing information is not archived.

Full-text Indexing (Keyword Indexing)

For full-text indexing, each event (CEF or non-CEF) received on Logger is scanned and divided into keywords and stored on the Logger. The full-text search options control the manner in which an event is tokenized as described in [“Full-text Search Options” on page 322](#).

Full-text indexing is automatically enabled at Logger initialization time. You cannot disable it. For details about enabling full-text indexing, see [“Enabling Indexing” on page 124](#).

Field-based Indexing

The field-based indexing capability allows for fields of events to be indexed. The fields are based on a predetermined schema. The Logger’s reports and the field search method utilize these indexed fields to yield significant search and reporting performance gains.

Field-based indexing for a recommended set of fields is automatically enabled at Logger initialization time. You can add more fields to an index at any time. Once a field has been added, you cannot remove it.

A list of the default index fields, along with their field descriptions is available from the Logger Configuration menu. For instructions on how to view the default Logger Schema fields, see [“Viewing Default Fields” on page 324](#).

**Note**

- HP strongly recommends that you index fields that you will be using in search and report queries.
 - The `requestUrl` field is available for search and report queries; however, this field cannot be indexed.
 - The fields created when a predefined or user-defined rex parser parses the non-CEF events cannot be indexed using the field-based indexing capability. See [“Parsers” on page 278](#) for more information about rex parsers.
-

In addition to indexing the fields included in the field-based indexing list, Logger indexes event metadata fields—event time, Logger receipt time, and device address—for every event. The event metadata fields are also known as the “internal” fields and are in addition to the fields you can add through the Logger’s user interface.

The following fields are available for indexing. The fields that Logger starts indexing automatically after Logger initialization are indicated in **bold** font.

In addition to the following fields, the `requestUrl` field is available for search queries. However, this field **cannot** be indexed.

Index Fields		
<code>agentAddress</code>	<code>deviceCustomDate2</code>	<code>flexDate1</code>
<code>agentHostName</code>	<code>deviceCustomDate2Label</code>	<code>flexDate1Label</code>
<code>agentNtDomain</code>	<code>deviceCustomNumber1</code>	<code>filePath</code>
<code>agentSeverity</code>	<code>deviceCustomNumber1Label</code>	<code>flexNumber1</code>
<code>agentType</code>	<code>deviceCustomNumber2</code>	<code>flexNumber1Label</code>
<code>agentZone</code>	<code>deviceCustomNumber2Label</code>	<code>flexNumber2</code>
<code>agentZoneName</code>	<code>deviceCustomNumber3</code>	<code>flexNumber2Label</code>
<code>agentZoneResource</code>	<code>deviceCustomNumber3Label</code>	<code>flexString1</code>
<code>agentZoneURI</code>	<code>deviceCustomString1</code>	<code>flexString1Label</code>
<code>applicationProtocol</code>	<code>deviceCustomString1Label</code>	<code>flexString2</code>
<code>baseEventCount</code>	<code>deviceCustomString2</code>	<code>flexString2Label</code>
<code>bytesIn</code>	<code>deviceCustomString2Label</code>	<code>message</code>
<code>bytesOut</code>	<code>deviceCustomString3</code>	<code>name</code>
<code>categoryBehavior</code>	<code>deviceCustomString3Label</code>	<code>priority</code>
<code>categoryDeviceGroup</code>	<code>deviceCustomString4</code>	<code>requestClientApplication</code>
<code>categoryObject</code>	<code>deviceCustomString4Label</code>	<code>requestContext</code>
<code>categoryOutcome</code>	<code>deviceCustomString5</code>	<code>requestMethod</code>
<code>categorySignificance</code>	<code>deviceCustomString5Label</code>	<code>requestUrlFilename</code>
<code>categoryTechnique</code>	<code>deviceCustomString6</code>	<code>requestUrlQuery</code>
<code>customerName</code>	<code>deviceCustomString6Label</code>	<code>sessionId</code>
<code>destinationAddress</code>	<code>deviceEventCategory</code>	<code>sourceAddress</code>
<code>destinationDnsDomain</code>	<code>deviceEventClassId</code>	<code>sourceHostName</code>
<code>destinationHostName</code>	<code>deviceExternalId</code>	<code>sourceMacAddress</code>
<code>destinationMacAddress</code>	<code>deviceHostName</code>	<code>sourceNtDomain</code>
<code>destinationNtDomain</code>	<code>deviceInboundInterface</code>	<code>sourcePort</code>
<code>destinationPort</code>	<code>deviceOutboundInterface</code>	<code>sourceProcessName</code>
<code>destinationProcessName</code>	<code>deviceProduct</code>	<code>sourceServiceName</code>
<code>destinationServiceName</code>	<code>deviceReceiptTime</code>	<code>sourceTranslatedAddress</code>
<code>destinationTranslatedAddress</code>	<code>deviceSeverity</code>	<code>sourceUserId</code>
<code>destinationUserPrivileges</code>	<code>deviceVendor</code>	<code>sourceUserName</code>
<code>destinationUserId</code>	<code>deviceVersion</code>	<code>sourceUserPrivileges</code>
<code>destinationUserName</code>	<code>deviceZone</code>	<code>sourceZone</code>

Index Fields		
destinationZone	deviceZoneName	sourceZoneName
destinationZoneName	deviceZoneResource	sourcezoneResource
destinationZoneResource	deviceZoneURI	sourceZoneURI
destinationZoneURI	endTime	startTime
deviceAction	eventId	transportProtocol
deviceAddress	externalId	type
deviceCustomDate1	fileName	vulnerabilityExternalID
deviceCustomDate1Label		VulnerabilityURI

Guidelines for Field-based Indexing

Make sure you are familiar with these guidelines before you index any fields:

- Events are indexed by the fields in the “Indexed fields” list (on the Search Indexes page) and the default event metadata fields—event time, Logger receipt time, and device address.
- You can index up to 123 fields on Logger. This number includes the custom schema fields you may have added to your Logger.
- Once a field has been added to the index, it cannot be unindexed.
- Only users belonging to a System Admin Group can add fields to index.
- After you add a field to the index, Logger might not immediately start indexing on that field. Therefore, allow some time between adding a field and using it in the search query. If Logger is in the process of indexing on a field and you use that field to run a search query, the search performance for that operation will be slower than expected.
- If an event field contains data of unexpected type (for example, a string when an integer is expected), the data is ignored. Therefore, search for that data value will not yield any results. For example, if the port field contains a value 8080A (alphanumeric) instead of 8080 (numeric), the alphanumeric value is ignored.
- For faster report generation, ALL fields of a report (including the fields being displayed in the report) need to be indexed. That is, in addition to the fields in the WHERE clause of the query, the fields in the SELECT clause also need to be indexed.
- For optimal search performance, make sure that event fields on ALL peers are indexed for the time range specified in a query. If an event field is indexed on a Logger but not on its peers for a specific time range, a distributed search will run slower on the Loggers. However, it will run at optimal speed on the local Logger. Therefore, the search performance in such a setup will be slow.
- Although the `requestUrl` field is available for search and report queries, it cannot be indexed. Including this field in such queries will result in the query running slower than a search performed on indexed data.

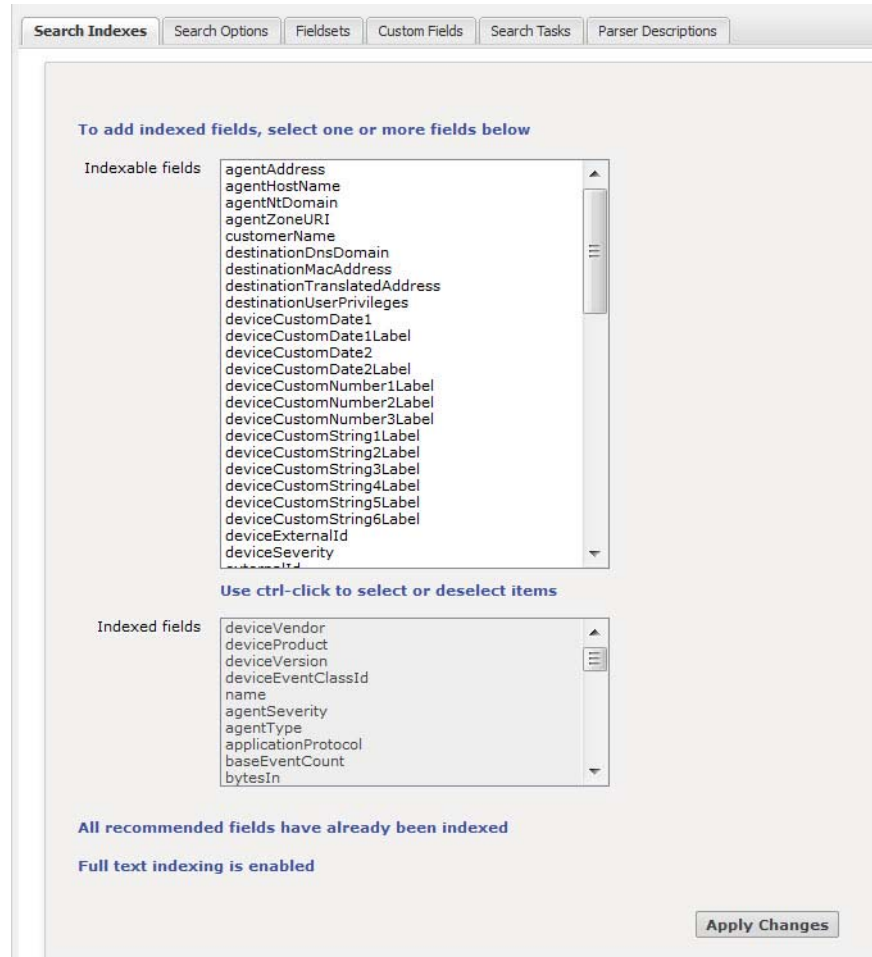
Enabling Indexing

Indexing is automatically enabled when Logger is initialized. You cannot disable indexing, however, you can add fields to the field-based indexing at any time.

Adding Fields to Field-based Index

To add fields to the field-based index:

- 1 Click **Configuration** from the top-level menu bar.
- 2 Click **Search**.
- 3 In the **Search Indexes** tab, select the fields from the Indexable Fields list.



- 4 To select multiple fields at the same time, hold the Ctrl key down and click on the fields.
- 5 Click **Apply Changes**.

Super Indexing

In addition to full text and field based indexing, Logger 5.5 and later creates super indexes for common IP address, host name, and user name fields. Super indexes enable Logger to quickly determine whether a particular field value has been stored on this Logger, and if it has, to narrow down the search to sections of data where that field value exists. Therefore, searches that can take advantage of super indexes return very quickly if there are no hits and return results more quickly than regular searches when there are very few hits.

- For information on how to use super indexes, see [“Searching for Rare Field Values” on page 83](#).
- A complete list of super-indexed fields is included in [Table 4-2, “Fields With Super Indexes,” on page 84](#).

Saving Queries (Saved Filters and Searches)

If you need to run the same search query regularly, you can save it in as a filter or as a saved search.

- Saving it as a filter saves the query expression, but does not save the time range or the fieldset information.
- Saving it as a saved search saves the query expression and the time range that you specified.

For information about Saved Search Alerts, see [“Creating and Managing Saved Search Alerts” on page 302](#).

Saving a Query

To save a query:

- 1 Define a query as described in [“Searching for Events on Logger” on page 106](#) or [“Using the Advanced Search Builder Tool” on page 96](#).
- 2 Click the Save icon (📌) and enter a name for the query in the Name field, as shown in the following figure.

- 3 In the Save as field, select whether you want to save this query as a filter, as a saved search, or as a Dashboard panel.
 - ◆ If you select to **save as a Saved Search**, you can either keep the saved query as Saved Search or change it to a Scheduled Alert by specifying a schedule based on which the query runs periodically and generates alerts.
 - ◆ If you choose to **schedule the Saved Search**, you can either specify the schedule in the following screens or skip it for now.
 - ◆ If the search query includes an aggregation operator such as chart or top, a third option to save the query for a **Dashboard panel** is also displayed.

If you select this option, you need to enter the following parameters.

Parameter	Description
Title	Enter a meaningful name for the panel that will be added to the Dashboard.
Saved search	Select an existing saved search from the drop-down box that will be overwritten with this query. OR Select "New saved search" to create a new saved search query. Enter the new name in the text box.
Dashboard	Select an existing Dashboard from the drop-down box to which the Search Results panel will be added. OR Select "New dashboard" to add the Search Results panel to a new Dashboard. Enter the name of the new Dashboard in the "Dashboard Name" field.
Panel type	Select the type of panel: <ul style="list-style-type: none"> Chart—Displays search results in a chart form Table—Displays search results in a table form Chart and Table—Adds two panels, one for displaying search results in the chart form and the other for displaying search results in the table form
Chart type	Type of chart to display matching events. You can select from: Column, Bar, Pie, Area, Line, Stacked Column, Stacked Bar. Default: Column
Chart limit	<i>Only applicable to Search Result Chart panels.</i> Number of unique values to plot. Default: 10

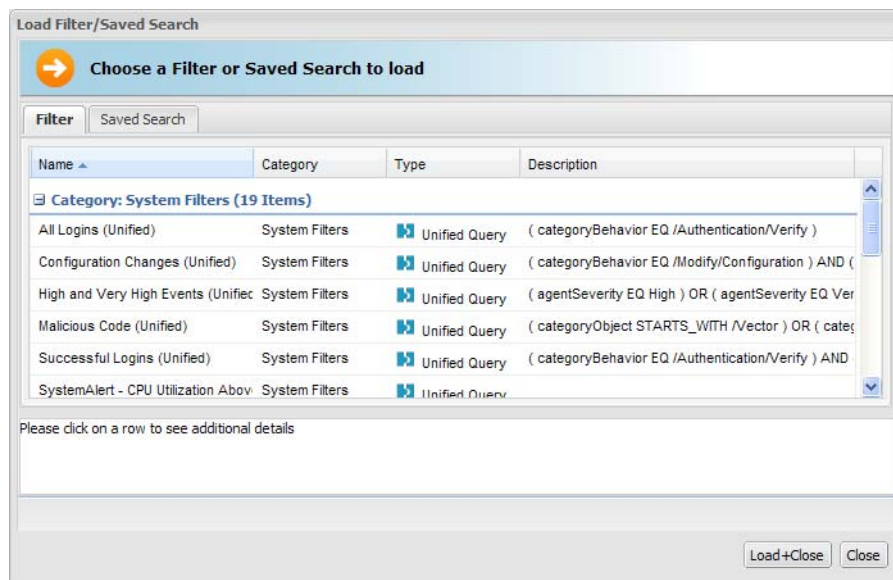
- 4 Click **Save**.

Using a Saved Filter or a Saved Search

The Load Search Filter/Saved Search interface enables you to quickly locate system filters, search filters, and saved searches. Your system provides pre-defined search filters that you can select to run. These are explained in [“System Filters/Predefined Filters” on page 128](#).

To use a saved filter (or a saved search):

- 1 Click **Analyze > Search**.
- 2 Click the Load a Saved Filter icon (🔍) to view a list of all the saved filters and saved searches to display the Load Filter/Saved Search interface, as shown in the following figure.



The Load Filter/Saved Search interface enables you to quickly locate the saved filters and the saved search queries. Click on any of the column names to sort information. To view details of a filter or a saved search, click its row. Details are displayed in the text box below.

To reload a filter, select the filter or saved search you want to use and click **Load+Close**. The filter rows display the search query.

To reload a saved query, click the **Saved Searches** tab, select a search, and click **Load+Close**.

System Filters/Predefined Filters

Your Logger ships with a number of predefined filters, also known as system filters. These filters define queries for commonly searched events. For example, unsuccessful login attempts or the number of events by source. Filter queries are available as Unified queries

and as Regular Expression queries. Unified queries can be used for searching and reporting while Regular Expression queries are for defining alerts and forwarders.



- Even though the Category - System Alert filters (listed in the last section of the following table) are displayed on the user interface of the software version of Logger, these filters do not apply to it.
- To effectively use the Firewall or UNIX Server use case filters (listed in the following table), define device groups that include the firewall devices or UNIX servers that you are interested in and then constrain your search to those device groups. If you do not create device groups specific to device types, the search results would match all Deny, Drop, or Permit events from all devices instead of only the firewall devices. Similarly, the "Unix-IO Errors and Warnings" filter would include IO errors and warnings from all devices and not only the UNIX servers.

The following is a list of all the system filters. For a description of each filter, see ["System Filters" on page 631](#). To use a predefined system filter, follow instructions in ["Using a Saved Filter or a Saved Search" on page 128](#).

Table 4-7 System Filters

Category	Unified Query Filters	Regular Expression Query Filters
Login Status use case	All Logins	All Logins (Non-CEF) All Logins (CEF format)
	Unsuccessful Logins	Unsuccessful Logins (Non-CEF) Unsuccessful Logins (CEF format)
	Successful Logins	Successful Logins (Non-CEF) Successful Logins (CEF format)
	Failed Logins	
Configuration	Configuration Changes	System configuration changes (CEF format)
Events use case	High and Very High Severity Events	High and Very High Severity CEF events
	Event Counts by Source	
	Event Counts by Destination	
		All CEF events
Intrusion use case	Malicious Code	Malicious Code (CEF format)
Firewall use case	Deny (Firewall Deny)	
	Drop (Firewall Drop)	
	Permit (Firewall Permit)	
Network use case	DHCP Lease Events	
	Port Links Up and Down	
	Protocol Links Up and Down	
Connector System Status use case	CPU Utilization by Connector Host	

Table 4-7 System Filters (Continued)

Category	Unified Query Filters	Regular Expression Query Filters
UNIX Server use case	Disk Utilization by Connector Host	
	Memory Utilization by Connector Host	
	CRON related events	
	IO Errors and Warnings	
	PAM and Sudo Messages	
	Password Changes	
	SAMBA Events	
	SSH Authentications	
	User and Group Additions	
	User and Group Deletions	
Windows Events use case	Account Added to Global Group	
	Account Added to Global Group (CEF)	
	Audit Policy Change	
	Audit Policy Change (CEF)	
	Change Password Attempt	
	Change Password Attempt (CEF)	
	Global Group Created	
	Global Group Created (CEF)	
	Logon Bad User Name or Password	
	Logon Bad User Name or Password (CEF)	
	Logon Local User	
	Logon Local User (CEF)	
	Logon Remote User	
	Logon Remote User (CEF)	
	Logon Unexpected Failure	
	Logon Unexpected Failure (CEF)	
	New Process Creation	
	New Process Creation (CEF)	
	Pre-Authentication Failure	
	Pre-Authentication Failure (CEF)	
	Special Privileges Assigned to New Logon	
	Special Privileges Assigned to New Logon (CEF)	

Table 4-7 System Filters (Continued)

Category	Unified Query Filters	Regular Expression Query Filters
System Alerts	User Account Changed	
	User Account Changed (CEF)	
	User Account Password Set	
	User Account Password Set (CEF)	
	Windows Events (CEF)	
	<p>The following filters search for specific internal alert events, which are written in CEF format to a special Internal Storage Group. These filters are available for both search methods. In addition to the following filters, you can define your own alerts based on the system health events listed in “System Health Events” on page 408.</p> <p>Note: Although these filters are displayed on the software version of Logger, these do not apply to it.</p>	
	CPU Utilization Above 90 Percent	CPU Utilization Above 90 Percent
	CPU Utilization Above 95 Percent	CPU Utilization Above 95 Percent
	Disk Failure	Disk Failure
	Root Partition Below 10 Percent	Root Partition Below 10 Percent
	Root Partition Below 5 Percent	Root Partition Below 5 Percent
	Device Configuration Changes	Device Configuration Changes
	Filter Configuration Changes	Filter Configuration Changes
	High CPU Temperature	High CPU Temperature
		Bad Fan
	Power Supply Failure	Power Supply Failure
	RAID Controller Issue	RAID Controller Issue
	RAID Status Battery Failure	RAID Status Battery Failure
	RAID Status Disk Failure	RAID Status Disk Failure
	Storage Configuration Changes	Storage Configuration Changes
	Storage Group Usage Above 90%	Storage Group Usage Above 90%
	Storage Group Usage Above 95%	Storage Group Usage Above 95%
	Zero Events Incoming	Zero Events Incoming
	Zero Events Outgoing	Zero Events Outgoing

Using a System Filter

To use a predefined system filter, follow instructions in [“Using a Saved Filter or a Saved Search” on page 128](#).

Alerts

You can configure your Logger to alert you by e-mail, an SNMP trap, or a Syslog message when a new event that matches a specific query is received or when a specified number of matches occur within a given time threshold.

Only regular expressions can be used in queries specified for alerts.

Audit events for alerts are only written to the Internal Storage Group and not forwarded to ESM by default. If you need to forward these audit events to ESM, please contact customer support for assistance.



Note

This change only applies to audit events generated for alerts; other audit events are unaffected.

Viewing Alerts

In addition to receiving an alert through the methods mentioned above, you can also view them through the user interface.

The Alert sub-tab under the Analyze tab presents a user interface that is similar to Search. From this page, you view Alerts and the base events that triggered them, as shown in the following figure.

When you create Alerts (see [“Alerts” on page 296](#)), you name them, and you can choose to view only events associated with a particular Alert. The default is All Alerts.

To view Alerts, choose a predefined time range, such as “Last 2 hours” or “Today,” or choose “Custom Time Range” to reveal additional fields for specifying a time range manually. This aspect works like Search. Refer to [“Time Range” on page 86](#) for more detail.

Receiving Alerts for Events

To receive alerts:

- 1 Configure the Logger’s SMTP with the desired e-mail address destination (see [“Static Routes” on page 361](#)) or create an SNMP Destination (see [“Sending Notifications to SNMP Destinations” on page 307](#)) or Syslog Destination (see [“Sending Notifications to Syslog Destinations” on page 308](#)).



Note

Number of destinations per alert:

- E-mail: Multiple, each separated by a comma.
 - SNMP: One
 - Syslog: One
-

- 2 Create a query to find the events of interest; save the query as a filter. (See [“Saving Queries \(Saved Filters and Searches\)” on page 126](#).)

- 3 Create an Alert that uses the new filter and specify match count and threshold (see “Alerts” on page 296.) Enable the new Alert.

Summary Analyze Dashboards Reports Configuration System Admin admin Logout

Show: All Alerts Within: Last 2 hours

Base Event Fields: All Fields Go! Export Results Auto Update: 5 min Paused

Alerts: 25 Status: Paused

Page Start: 14/May/2008 13:11:38 -0700 Page End: 14/May/2008 13:11:39 -0700

Next >

Time (Event Time)	Alert Name	Base Event Count	Time Threshold	Matched Events
14/May/2008 13:11:38 -0700	Email Alert	1	2	1

Base Event (1 found)

Time (Event Time)	Device	Logger	Name	Severity	Receipt Time	Device Vendor	Device Product
14/May/2008 13:07:00 -0700	127.0.0.1	Local	Logger Internal Event	1	14/May/2008 13:11:38 -0700	ArcSight	Logger

14/May/2008 13:11:38 -0700 Email Alert 1 2 1

Base Event (1 found)

Time (Event Time)	Device	Logger	Name	Severity	Receipt Time	Device Vendor	Device Product
14/May/2008 13:07:00 -0700	127.0.0.1	Local	Logger Internal Event	1	14/May/2008 13:11:38 -0700	ArcSight	Logger

14/May/2008 13:11:38 -0700 Email Alert 1 2 1

Base Event (1 found)

Time (Event Time)	Device	Logger	Name	Severity	Receipt Time	Device Vendor	Device Product
14/May/2008 13:07:00 -0700	127.0.0.1	Local	Logger Internal Event	1	14/May/2008 13:11:38 -0700	ArcSight	Logger

14/May/2008 13:11:38 -0700 Email Alert 1 2 1

Base Event (1 found)

Time (Event Time)	Device	Logger	Name	Severity	Receipt Time	Device Vendor	Device Product
14/May/2008 13:07:00 -0700	127.0.0.1	Local	Logger Internal Event	1	14/May/2008 13:11:38 -0700	ArcSight	Logger

14/May/2008 13:11:38 -0700 Email Alert 1 2 1

Base Event (1 found)

Time (Event Time)	Device	Logger	Name	Severity	Receipt Time	Device Vendor	Device Product
14/May/2008 13:07:00 -0700	127.0.0.1	Local	Logger Internal Event	1	14/May/2008 13:11:38 -0700	ArcSight	Logger

Base Event Fields

Events that are labeled 'Action Engine' are Alert events. Other events are base events--that is, the events that triggered the Alert.

Go, Export, and Auto Update Options

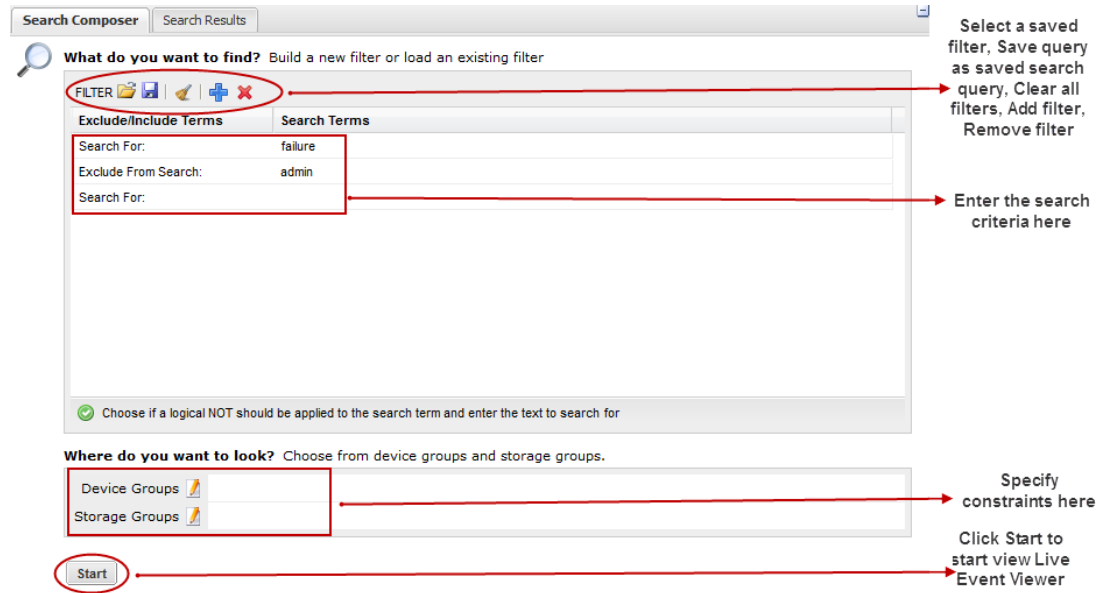
The **Go** and **Export Results** buttons and the **Auto Update** option accomplish the same tasks in both the Search and Alert pages. For more information, see “Searching for Events on Logger” on page 106, “Understanding the Search Results Display” on page 109, “Viewing Alerts” on page 132, and “Advanced Search Options” on page 107.

Live Event Viewer

The Live Event Viewer provides real-time view of the incoming events that match the criteria you specify. This functionality is useful in environments where the need to view an event quickly is important; for example, a financial institution might be interested in viewing a specific transaction type as soon as it occurs. Because the latency between the events arriving at Logger and the display time is quite less, events might not have been indexed on Logger before being displayed.

The Live Event Viewer composes of two tabs—Search Composer and Search Results. The Search Composer is for defining the search criteria and the Search Results tab displays the matching events in real time.

The following figure shows the Search Composer. If more than one filter is specified, the resulting query uses the AND operator to combine them. For example, if the first filter searches for “failure” and the second filter **excludes** “admin”, the resulting query is “failure AND NOT admin”.




The Search Results tab provides the Play, Pause, Stop, Clear, and Export buttons that enable you to control the display in a manner similar to any electronic device, as shown in the following figure.



The following list highlights the features of Search Results display:


- Events are displayed in the raw event format and not in the columnar, table form as displayed in the Search Results page (**Analyze > Search**) when you run a search query.

- A user can launch a maximum of one Live Event Viewer. There can be a maximum of five Live Event Viewers running on Logger at any time.
- The regular expression search method is used to identify matching events. Therefore, you can specify regular expressions as the search term in the Search Composer.
- Buffer Size defines the maximum number of events displayed in the Viewer. By default, the Buffer Size is 1000, however, it can be set to any number between the range of 20 and 5000.
- By default, the search is run for 15 minutes and then stopped to preserve system resources. If you need to run the search for longer than 15 minutes, click the  icon next to the countdown timer to reset the timer to 15 minutes.
- When you click Pause, the Search Results display is frozen. However, the search operation continues in the background and the new matching events are buffered until a maximum of 1000 events have been buffered or the search timer, which continues to count down even when the Search Results display is frozen, reaches 00:00.
- If the timer has not reached 00:00, you can click Play to resume the paused search operation. When you click Play, the buffered events are displayed. The newly found events are appended to the previously found events on the Search Results display screen.
- When you click Stop, the search for matching events and the countdown of the search timer stop. When you click Play, the search is started afresh—the currently displayed events are cleared from the Search Results screen, the search timer is reset to 15 minutes, and the search starts again.
- You must stop the search operation to export the matching events.

To launch a Live Event Viewer:





Live Event Viewer is a resource-intensive application that can impact the overall performance of your Logger if run for a long period of time. Therefore, use this feature selectively and for short periods of time.


- 1 Click **Analyze > Live Event Viewer** from the top-level menu bar.
- 2 In the Search Composer tab, enter the search terms or click the () icon to select a saved filter.
- 3 You can enter search terms that the event must contain (Search For:) or terms that the events must not contain (Exclude From Search:). Click the "Search For:" field to display a drop-down list from which you can select "Exclude From Search:".

If more than one filter is specified, the AND operator is used to combine them in the resulting search query.

To add additional filters, click the () icon.

To remove a filter line, click the () icon.

To remove all filters, click the () icon.

- 4 Enter constraints to limit your search to specific device groups, devices, or storage groups in the "Where do you want to look?" section. Click the () icon to display a list from which you can choose the constraints.
- 5 Click **Start**.
- 6 The search results are automatically displayed in the Search Results display screen.

To update the Live Event Viewer query:

- 1 In the Search Composer tab of the Live Event Viewer, update the search terms.
- 2 Click **Stop** first, then **Start** to start search using the new search terms.

To export Search Results display:

- 1 Make sure you have stopped the Live Event Viewer. To do so, click the (■) icon in the Search Results display window.
- 2 Click the (⬇) icon to open the Export Options window.
- 3 Follow [Step 3](#) onward in “[To export search results:](#)” on [page 119](#) to export the displayed search results.

Chapter 5

Reporting

Reporting is an essential tool for communicating the state of your network security to internal and external stakeholders. Reports are captured views or summaries of events that you can view from the Logger Reports tab or export for sharing in a variety of file formats.



Reporting is not available for trial Loggers. To take advantage of this feature, you need the Enterprise version.

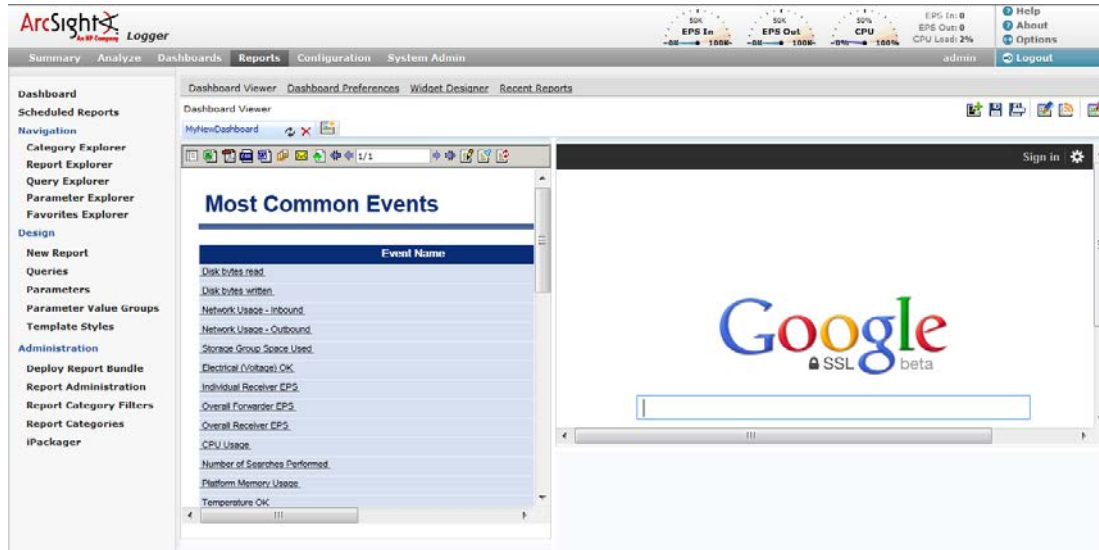
This chapter includes information on the following areas of Logger reporting:

- ["The Reports Home Page" on page 137](#)
- ["Explorers" on page 139](#)
- ["Categories" on page 147](#)
- ["Dashboards" on page 153](#)
- ["Widgets" on page 159](#)
- ["Using Dashboards Created in Pre-5.2 Logger" on page 162](#)
- ["Running, Viewing, and Publishing Reports" on page 172](#)
- ["Designing Reports" on page 185](#)
- ["Scheduling Reports" on page 233](#)
- ["Deploying a Report Bundle" on page 236](#)
- ["Report Server Administration" on page 238](#)
- ["Using Report Category Filters" on page 239](#)
- ["Backup and Restore of Report Content" on page 240](#)
- ["iPackager" on page 240](#)

The Reports Home Page

To get to the Reports home page from elsewhere in the Logger UI, click the **Reports** link on top of the page. Alternatively, click **Dashboard** on the left panel menu from within Reports, to display the Reports Home page.

The Dashboard Viewer opens and if a dashboard is configured to display, the Reports Home page shows the selected **Dashboard** view.



If there is no Dashboard display configured and selected, the Reports home page shows an empty Dashboard Viewer.

On top of the Reports home page there are links for Dashboard Viewer, Dashboard Preferences, Widget Designer, Recent Reports, Classic Viewer, Classic Designer, and Classic Preferences.

The **Dashboard Viewer** page allows you to view your dashboards. On the upper right corner of the Dashboard Viewer page there are buttons that you can use to add a widget to the dashboard, save edit, and subscribe the dashboard.

The **Dashboard Preferences** link allows you to specify a default dashboard to display as your Reports home page and also to display multiple dashboards as tabs in the Reports home page.

The **Widget Designer** page allows you to create a widget displaying either a report or a web link. You can place the widget in the dashboard from the Dashboard Viewer page.

The **Recent Reports** link shows the report execution status that lists the status of currently running, recently run, or accessed reports. By default, all reports except the completed scheduled reports are displayed; however, you can restrict the list by defining filter criteria. The Execution Type column under the Report Execution Status drop down indicates whether the report was run in the background or was run using quick run. A report run on ad-hoc basis is listed on the Report Execution Status page for 24 hours; however reports run in the background are listed for a longer period of time. You can select a report to run a report using different field values as well as re-run a report using the same values used in the original run.

To get started by creating a dashboard to show as your default Reports Home page, see [“Dashboards” on page 153](#) and [“Designing Dashboards” on page 154](#).

Dashboards created in pre-5.2 Logger can be viewed and edited from the **Classic Viewer**, **Classic Designer**, and **Classic Preferences** links. See [“Using Dashboards Created in Pre-5.2 Logger” on page 162](#) for information on how to view or edit dashboards created in pre-5.2 Loggers.



Do **not** use these links to create new dashboards. Use the Dashboard Viewer link to create new Dashboards. See [“Designing Dashboards” on page 154](#) for details.

To get started by running and viewing reports, see [“Running, Viewing, and Publishing Reports” on page 172](#). You can run, publish, and save the results of any type of report. For information on these common reporting tasks available on all reports, see [“Task Options on Available Reports” on page 173](#).)

You can schedule any report to run once at a later date or on a specified frequency (such as daily or weekly). Monthly reports cannot be scheduled currently. For more on this, see [“Scheduling Reports” on page 233](#).

The left pane of the Reports home page provides links that make it easy to navigate to a report, design a new report and links for report administration.

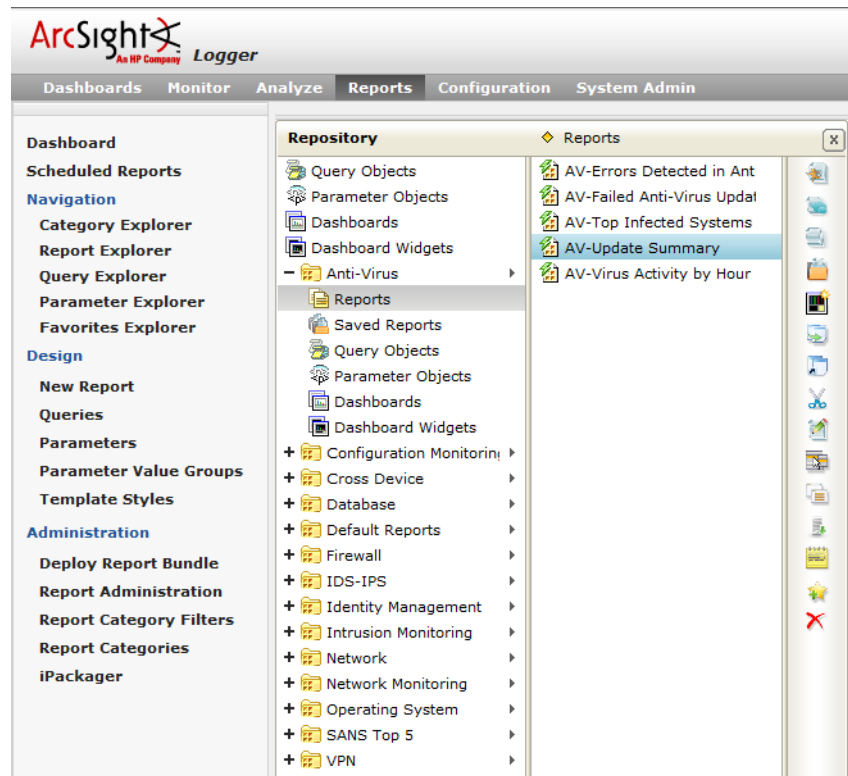
The various explorer links listed under **Navigation** in the left panel allow you to navigate to a report, query, or parameter. The links listed under **Design** allow you to create a new report and create, view and configure any object of the Reporting feature. The links listed under **Administration** allow you to do administrative tasks with reports.

Explorers

Use the explorers listed under the **Navigation** heading in the left pane to navigate to a desired report, query, parameter, dashboard, or dashboard widget that already exists in Logger. You can run a report, publish it, and customize it from the Report explorer. You can create and edit a query or parameter object from its respective explorer. You can select multiple objects in any of the explorers and perform actions such as set access rights on all of them, copy, cut, paste or delete actions on all of them. The following explorers are available.

Category Explorer

Reports and report objects, such as queries and parameters can be organized and grouped based on their function. Such functional groups are called Categories. For example, a report pertaining to a database can be stored under the Database category. The Category Explorer brings together all categorized reports and report objects in one central location.



The Category Explorer comes with some pre-defined, commonly used categories. You can add custom categories based on your requirements. You can use the Category Explorer to get an overview of all reports, published reports only, queries, parameters, dashboards, and dashboard widgets that were saved in a particular category.

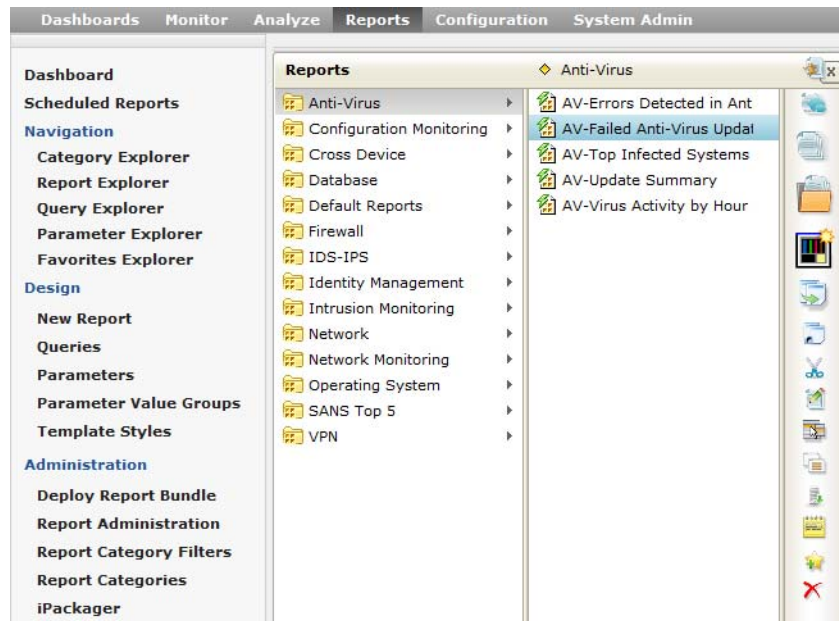
When you click the Category Explorer link, a Repository pane appears to the right of the navigation pane listing all the available categories. Navigate to a report or another object by clicking on the category to expand it. Click on any object stored in that category and another pane will open up to the right of it listing the contents stored under it. For example, to get to a report from the Category Explorer, click the Category Explorer link, click the category link in the Repository pane to expand it, click Reports under the category, and click on the report of your choice and select an action to be performed (such as run report) on the report by clicking on one of the buttons in the right menu bar.

See [“List of Buttons in the Explorers” on page 144](#) for an explanation of each available button.

For a complete list and description of reports available on Logger, see [Appendix F, Logger Content, on page 593](#).

Report Explorer

The Report Explorer is the central location for viewing (publishing), running, or editing a report. However, you cannot create a new report from the Report Explorer.




To create a new report, click the **New Report** link under the Design section in the left pane. The Report Explorer lists the various pre-defined categories. Clicking on any of the category opens another pane to its right which displays all the reports that have been stored in that category.

See [“List of Buttons in the Explorers” on page 144](#) for an explanation of each available button.

Query Explorer

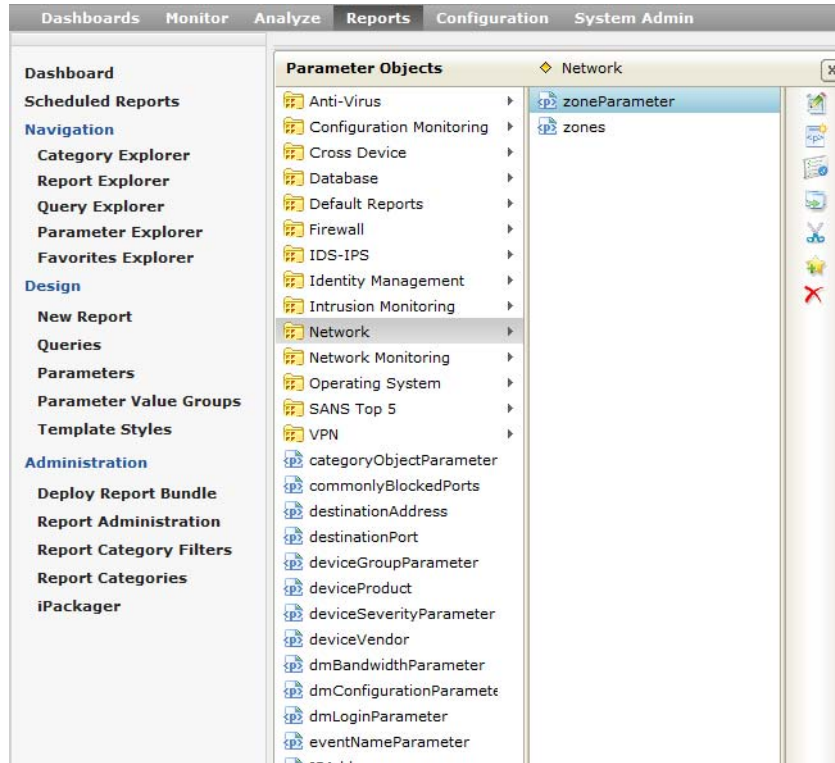
Queries are made up of query objects or parameters. These query objects can be grouped in categories as well. When you click on the Query Explorer link, it opens a list of some pre-defined query objects that have not been categorized and the various pre-defined categories under which you can store any newly created objects.



You can use the Query Explorer to create new queries, view the properties for saved queries and edit the saved queries' properties. See ["List of Buttons in the Explorers"](#) on [page 144](#) for an explanation of each available button.

Parameter Explorer

If you click the Parameter Explorer link, it opens a panel that lists the predefined categories under which you can store newly created parameters and some pre-defined parameters that have not been categorized.



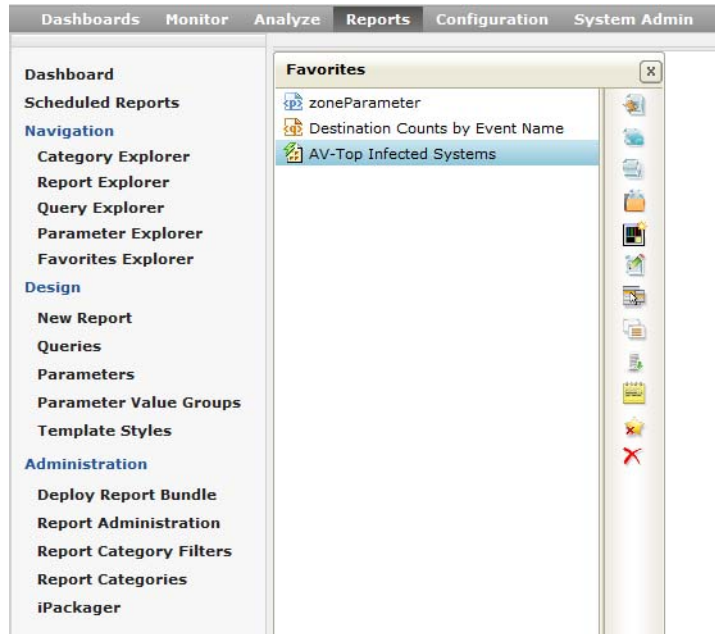
You can use the Parameter Explorer to create new parameters, view the properties for existing parameters and edit the existing parameters' properties.

See [List of Buttons in the Explorers](#) for an explanation of each button available from the explorer.

For a complete list of parameters available on Logger, see [“Parameters” on page 626](#).

Favorites Explorer

For quick access, you can mark any report, query, parameter, dashboard, or dashboard widget as a favorite. Anything you mark as a favorite will be listed in the Favorites Explorer. However, objects listed in the Favorites Explorer cannot be organized into categories.





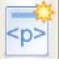









Any action that you can perform on an object from its own explorer can also be done to it from the Favorites Explorer too. For example, a report can be run, saved, and published from the Report Explorer. If you have marked the report as a Favorite, you can run, save, and publish it from the Favorites Explorer too.











See [“List of Buttons in the Explorers” on page 144](#) for an explanation of each available button.

List of Buttons in the Explorers

The following table lists all buttons that are available in the button bars in the various explorers. Not all buttons will be available from all explorers. The buttons will appear only when you click on the object. For instance, if you click on a report in the Reports pane of an explorer, all report-related buttons will appear in the button bar to the right of the Reports pane.

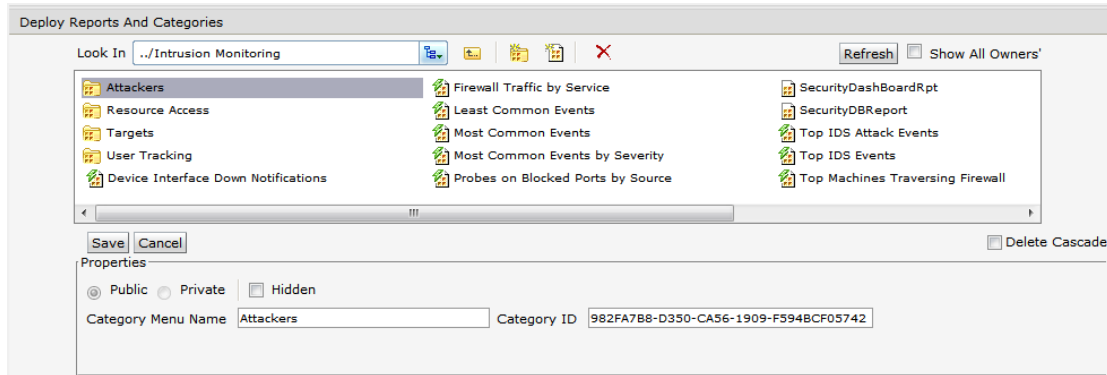
Buttons	Available in Explorer...	Used for
	All Explorers	Add to Favorites Any category, report, query, or parameter can be marked as a favorite. The Favorites Explorer lists all objects marked as favorites.
	Favorites Explorer	Delete From Favorites Any object that appears in the Favorites folder can be removed from there by clicking this button.

Buttons	Available in Explorer...	Used for
	Category Explorer Parameter Explorer	Create Parameter Object A new parameter can be created by clicking on this button. New parameters can be created from within the Category Explorer or the Parameter Explorer.
	Category Explorer Query Explorer	Create Query Object A new query can be created by clicking on this button. New queries can be created from within the Category Explorer or the Query Explorer.
	All Explorers	Delete Use this button to delete an object.
	All Explorers	Refresh Use this button to refresh an Explorer page.
	All Explorers	Properties Use this button to view object properties such as report properties, query properties, or parameter properties.
	Category Explorer Parameter Explorer Query Explorer	Edit <Object> Details Object can be a query, parameter or a dashboard widget In the case of reports, it stands for Customize Reports
	All Explorers	Copy <Object> Object can be a report, query, parameter, dashboard, or a dashboard widget
	All Explorers	Cut <Object> Object can be a report, query, parameter, dashboard, or a dashboard widget
	Category Explorer Report Explorer Favorites Explorer	Quick Run with Default Options Runs the report using default data filtering configuration, which was set at report deploy time. Provides options to change start and end time parameters, storage groups, and devices included in the scope of the report run. See also "To run and view a report:" on page 175 and "Quick Run with Default Options / Run In Background Report Parameters" on page 175
	Category Explorer Report Explorer Favorites Explorer	Run in Background Use this option to run reports that take long time to generate or the ones that are not required online immediately. See also "To run and view a report:" on page 175 and "Quick Run with Default Options / Run In Background Report Parameters" on page 175

Buttons	Available in Explorer...	Used for
	Category Explorer Report Explorer Favorites Explorer	Run Report Provides options to modify the data filter criteria used by the report query for this run. You can specify a maximum number of rows to include in the report, and perform various comparison and logical operations on event fields. See also "To run and view a report:" on page 175 and "Run Report Parameters" on page 178 .
	Parameter Explorer Category Explorer Favorites Explorer	Parameter Value Groups
	Category Explorer Report Explorer Favorites Explorer	Open Listing Page This button opens a page that displays reports (Adhoc reports, standard report and linked reports). You can view list of reports under selected category and take actions on a selected report. The complete path where report resides is shown in Report Listing.
	Category Explorer Report Explorer Favorites Explorer	Copy Report as Link
	Category Explorer Report Explorer Favorites Explorer	Customize Report
	Category Explorer Report Explorer Favorites Explorer	Download Report
	Category Explorer Report Explorer Favorites Explorer	List Published Output Click List Published Output to view the list of published outputs of this report. The list page also displays the user name who generated (published) the report, time and expiry time of the report. From this page, you can view report output as well as user comments on the report.
	Category Explorer Report Explorer Favorites Explorer	View Descriptions
	Category Explorer Report Explorer Favorites Explorer	Create Dashboard Widget
	Category Explorer Favorites Explorer	View Dashboard

Categories

Reports, queries, and parameters can be organized and stored under categories for ease of access. You can create your own categories or edit the existing categories' properties by clicking on the Report Categories link in the left pane of the Reports Home page.



System Defined Categories

The following categories are system-defined (pre-defined) and ready for use. They are based on common areas of use.

- Anti-Virus
- Configuration Monitoring
- Cross Device
- Database
- Default Reports
- Firewall
- Identity Management
- IDS-IPS
- Intrusion Monitoring
- Network
- Network Monitoring
- Operating System
- SANS Top 5
- VPN

Anti-Virus

Use this category to store reports, queries, parameters, dashboards, and dashboard widgets that provide information on anti-virus activity, such as the anti-virus update status, virus activity by hour, and top infected systems.

For a complete list of reports, click the Anti-Virus category under Reports column in the **Report Explorer**.

Configuration Monitoring

Logger provides reports that address configuration monitoring. To access these reports, click the **Configuration Monitoring** category in the Reports column in the Reports Explorer.

For information how to run, view, and publish these reports, see [“Running, Viewing, and Publishing Reports” on page 172](#).

Cross Device Reports

These reports provide information on functions that apply to multiple kinds of devices, such as failed login attempts, bandwidth usage by hosts, and accounts created by user, and so on.

For a complete list of reports, click the **Cross Device** category under the Reports column in the Reports Explorer.

Database Reports

These reports provide information on database errors and warnings.

Default Reports

The user-generated reports are placed in this category.

Firewall Reports

These reports provide information on firewall activity, such as denied connections by port, address, and hour.

Identity Management Reports

This report provides information on the number of connections per user as reported by the Identity Management devices in your network.

IDS-IPS Reports

These reports provides information on activity involving Intrusion Detection and Prevention Systems, such as alert count by device, port, severity, top alert destinations, worm infected systems, and so on.

Intrusion Monitoring Reports

Logger provides reports that address intrusion monitoring. To access these reports, click the Reports Explorer link and click the **Intrusion Monitoring** category in the Reports column.

For example, reports are provided to track password changes, firewall configuration events, firewall traffic, top attackers traversing firewalls, and so forth.

For information how to run, view, and publish these reports, see [“Running, Viewing, and Publishing Reports” on page 172](#).

Network Reports

These reports provide information on activity involving network infrastructure, including interface status, device errors, SNMP authentication failures, and so on.

Network Monitoring Reports

Network Monitoring reports describe activities on Virtual Private Networks:

- Top VPN Accesses by User
- Top VPN Event Destinations and Sources
- Top VPN Events
- VPN Connection Attempts
- VPN Connection Failures

Operating System Reports

These reports provide information on activity involving operating systems, such as login errors per user, user and user group creation and modification events.

SANS Top 5 Reports

Logger provides reports that address the “SANS Top 5 log reports” scenarios, all pre-built and available to run on-demand or schedule for a specified frequency. To access these reports, click the Report Explorer link in the left panel and click on **SANS Top 5** category under the Reports column.

For information on how to run, view, and publish these reports, see [“Running, Viewing, and Publishing Reports” on page 172](#).

The SANS Institute is a cooperative training, certification, and research organization with a focus on developing solutions for securing information against a variety of potential threats. SANS facilitates and supports a collaborative effort of a large number of security practitioners in various industries and sectors around the world to share experience, solutions, and resources related to information security. (“SANS” stands for “SysAdmin, Audit, Network, Security”; more information is available on their Web site at <http://www.sans.org/>.)

The “SANS Top 5” represents the current set of “most critical” log reports for a wide cross-section of the security community.

Here is a quote from the SANS Web site about the strategy and focus of the “SANS Top 5 Essential Log Reports”:

“The goal is to include reports that have the highest likelihood of identifying suspect activity, while generating the lowest number of false positive report entries. The log reports may not always clearly indicate the extent of an intrusion, but will at least give sufficient information to the appropriate administrator that suspect activity has been detected and requires further investigation.”



The SANS Top 5 list is meant to be reviewed on a regular basis. ArcSight can send updates for customers to deploy as new reports are required to meet new challenges presented by the dynamic threat-security environment in which networks are deployed.

The “SANS top 5” log reports cover the following five scenarios:

- Attempts to gain access through existing accounts
- Failed file or resource access attempts
- Unauthorized changes to users, groups and services

- Systems most vulnerable to attack
- Suspicious or unauthorized network traffic patterns

For a complete description of the SANS Top 5 log reports, see http://www.sans.org/resources/top5_logreports.pdf or look for associated topics in SANS “resources” on their Web site.

The Logger “SANS Top 5 Reports” offered to address these threat scenarios are:

- SANS Top 5 - 1 Number of Failed Logins
- SANS Top 5 - 1 Top Users with Failed Logins
- SANS Top 5 - 2 Failed Resource Access by Users and Drilldown
- SANS Top 5 - 2 Failed Resource Access Events and Drilldown
- SANS Top 5 - 3 Password Changes
- SANS Top 5 - 3 User Account Creations, Deletions, and Modifications
- SANS Top 5 - 4 Vulnerability Scanner Logs by Host or by Vulnerability
- SANS Top 5 - 5 Alerts from IDS
- SANS Top 5 - 5 IDS Signature Destinations and Source
- SANS Top 5 - 5 Top 10 Talkers
- SANS Top 5 - 5 Top 10 Types of Traffic
- SANS Top 5 - 5 Top Destination and Target IPs

VPN Reports

These reports provide information on activity involving VPN connections, including authentication errors, connection information such as counts, accepted and denied by address, and so on.



Tip

More Reports may be available for download as report packages on the HP Customer Support site (SSO). (For information about deploying report packages, see [“Deploying a Report Bundle” on page 236.](#))

Solution Reports

If any solution packages are installed on the Logger, they appear under this report group. Solution packages address specific compliance requirements or scenarios and are installed separately. Solutions Reports are available as add-on packages to Logger for specific compliance requirements or scenarios.



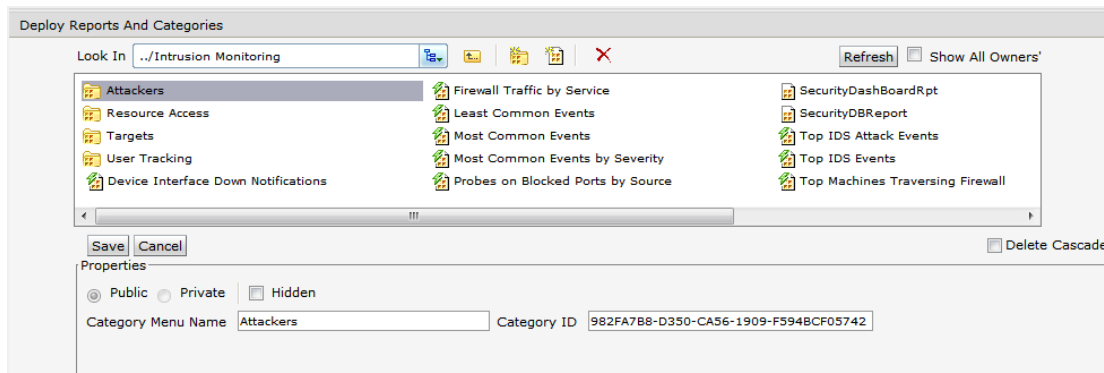
Tip

More Solution Reports may be available for download as report packages on the HP Customer Support site (SSO). (For information about deploying report packages, see [“Deploying a Report Bundle” on page 236.](#))


For information on deploying Solutions Packages, see [“Deploying a Report Bundle” on page 236.](#) Once deployed, these solution reports are listed in categories under the Solution Reports report group. To access these reports (once they are deployed), click Reports | Solutions Reports | *<report category name>* on the left menu, where *<report category name>* is the solution name, for example: PCI.

Adding a New Category

The Category Explorer comes with some system defined commonly used categories.



To add custom defined category:


- 1 Click the **Report Categories** link under the Administration heading in the left pane.
- 2 Click the **Add New Category**  button located on top of the box that displays all existing categories.
- 3 Define the properties for the new category and click the **Save** button.

Property	Used for...
Public	Setting this as Public makes the category available to everyone
Private	Setting this as Private make the category available to you only
Hidden	Check the Hidden checkbox, if you do not want to display this category in any of the dialogs and pages (except in the Reports Explorer). Mark a category as hidden to stop users from directly accessing the objects and reports stored in it.
Category Menu Name	Name of the Category
Category ID	Category ID should be unique across all the categories. By default, the Category ID is auto-generated by the system. To specify the Category ID manually, uncheck the System Generated checkbox and specify the category ID.
System Generated	To specify the Category ID manually, uncheck the System Generated checkbox and specify the category ID.
Delete Cascade	You can delete a category only if it is empty. To delete a category including its contents, check the Delete Cascade checkbox.



Note

Once set, Category ID and scope (Public / Private options) cannot be changed.


- 4 You can optionally add a report to the category. To do so, double-click any category to open it and click the **Add New Report**  button. Define the following properties in the Properties box:

Property	Used for...
Public	Setting this as Public makes the report available to everyone
Private	Setting this as Private make the report available to you only
Hidden	Check the Hidden checkbox, if you do not want to display this report in any of the dialogs and pages (except in the Reports Explorer). Mark a report as hidden to stop users from directly accessing it.
Report File	An existing data file from which a report is generated
Report Name	The Report Name has to be unique within a category
Report ID	A unique ID for the report that is auto-generated by the system by default when you run and publish the report. To manually enter an ID of your choice, uncheck the System Generated checkbox and enter an ID in the Report ID field.
Design Mode	Text in Design Mode indicates if the report was designed using Studio (Web Studio or Desktop Studio) or Adhoc Report Wizard.
Deployment Type	A report deployed as Read Only cannot be modified and uploaded with same name. A report deployed as Custom can be modified and uploaded with the same name.
Output Format	Output Formats in which this report can be generated. Formats not selected here will not be available for this report.
System Generated	To specify a Report ID manually, uncheck the System Generated checkbox and specify the Report ID.


Deleting an Existing Category

You can delete a category only if it is empty.

To delete an empty category:

Select the category in the **Category Explorer** and click the **Delete**  button in the vertical button bar to delete the selected category.

To delete a category including its contents:

Click the **Report Categories** link in the left pane, select the category, check the **Delete Cascade** checkbox and click the **Delete Selected Category**  button.




Note

If you attempt to delete a category that is not empty, and Delete Cascade checkbox is clear, a message "Failed to delete the category" will appear on top left of the page.

Placing a System Defined Query or Parameter into a Category


You can place a pre-defined query or parameter into a category. Use the cut/paste feature to do so because cutting and pasting it will preserve its ID.

To cut and paste a query/parameter:

- 1 Click the **Query Explorer** or **Parameter Explorer** link (depending on what you want to place in the category) in the left pane.
- 2 Click on the pre-defined query/parameter you want to move.
- 3 Click the **Cut Query Object/Cut Parameter Object**  button on the side button bar.
- 4 Click the category name under which you would like to place this query/parameter.



You cannot save a report in the root category. Save it in one of the existing subcategories, or save it in a new category.

- 5 Click the **Paste**  button on the side button bar.



Do not copy and paste a query or parameter to place it in a category. Doing so will give the query/parameter a new ID and render it unusable to reports or other existing objects that are using it.

You can schedule any report to run once at a later date or on a specified frequency (such as daily or weekly). Monthly reports cannot be scheduled currently. For more on this, see [“Scheduling Reports” on page 233](#).

You can run, publish, and save the results of any type of report. For information on these common reporting tasks available on all reports, see [“Running, Viewing, and Publishing Reports” on page 172](#) and [“Task Options on Available Reports” on page 173](#).

Dashboards

Dashboards display reporting data to provide a quick view of the latest information about network events. You can assemble various reports and external links onto a dashboard. However, you must place each report or link into its own widget and then place the widget in the dashboard. A dashboard can contain multiple widgets.

Placing reports on a dashboard gives you access to the most recently published results for those reports. Keep in mind, reports must be run and published in order for the results to be accessible on a dashboard viewer. If you schedule a report to run, publish, and save for a reasonable retention period (for example, one month), then those results will always be available for dashboard views.

For example, you can add one or more reports to a dashboard, and configure reports to *auto-refresh* (get results) on a specified interval (for example, every hour). The dashboard will access the latest published reports results, in this case, every hour. If you have also *scheduled* the reports to run and publish every hour, your dashboard will get up-to-date

results. This eliminates the need to manually run and view each report once per hour in order to get the same information updates.



Note

To view or edit Dashboards created in pre-5.2 Logger releases see [“Using Dashboards Created in Pre-5.2 Logger” on page 162](#).

The general steps to get started with using dashboards are as follows:

- 1** Create a new dashboard. See [“Creating a New Dashboard” on page 155](#) for details on this.
- 2** Create one widget for every report or web link you want to display on the dashboard. See [“Creating a New Widget” on page 159](#) for details on this.
- 3** Add the widgets to the dashboard. See [“Placing Widgets in a Dashboard” on page 162](#) for more details on this.
- 4** Optionally, you can configure the dashboard to appear as a tab in the Dashboard Viewer. See [“Viewing an Existing Dashboard in a Tab in the Dashboard Viewer” on page 156](#) for more details on this.

Viewing the Dashboard

The Dashboard Viewer is the home page for Reports. If no dashboard is configured and selected for display, the default Reports home page shows an empty Dashboard Viewer. If a dashboard is configured and selected for display, it is shown on the Dashboard Viewer page, and serves as the Reports Home page. If you are viewing other pages within the Reports tab, click **Dashboard** on the left panel to return to the Dashboard Viewer (Reports Home page).

Reports must be run and published first in order for the results to be accessible on a dashboard viewer. There are no options available to *run* reports from the Dashboard Viewer. You can only *view* saved or published reports.

You can set a dashboard to auto-refresh at a certain interval, but auto-refreshing a dashboard simply updates the dashboard display with the most recently published results; it does not run the report. We suggest using the dashboard refresh interval in conjunction with scheduled reports to ensure that report results are always published and retained long enough to be available to dashboards. For more information, see [“Designing Dashboards” on page 154](#) and [“Scheduling Reports” on page 233](#).

To run a report manually, click the Report Explorer in the left pane, click the category in which the report is stored, select the report and click the Run Report, Quick Run with Default Options, or Run in Background button. For more information on running and publishing reports, see [“Running, Viewing, and Publishing Reports” on page 172](#).

The Dashboard Viewer page displays the various items placed on the dashboard. If the dashboard includes reports, reports will show current data from recently run reports.

Designing Dashboards

Use the **Dashboard Viewer** page to create a new dashboard, name it, add items to it, and design the layout. You can design and save multiple dashboards, but only one at a time can be set as the default Dashboard Viewer for the Reports home page. Other dashboards can be saved for later use. Each dashboard can include multiple items (reports, use cases, and Web links).

To access the Dashboard Viewer, click the **Dashboard** link in the left pane.

What Items Can a Dashboard Include?


The following information is available for placement on a dashboard. However, each report or Web Link must be placed inside a widget and the widget in turn is placed into the dashboard. A dashboard can contain one or more widgets containing either of the following:


- **Reports**; any report can be included. The dashboard will show the latest published version of the report. Reports must be published in order for the report data to be accessible to users on the Dashboard View. If no published results are available for a report on a dashboard, the Dashboard View will display a message indicating this. When the report is published, a refresh of the Dashboard view will display the report.
- **External Links**; that is, any URL(s) that you want on-screen as a part of a particular Dashboard View

Creating a New Dashboard

The high-level steps to create a dashboard are described here. A detailed explanation of each of these steps is provided in the topics that follow. The Reports home page opens with the Dashboard Viewer open.

To add a new dashboard:

- 1 Click the **New Dashboard**  button located on the extreme right on any Dashboard tab.

This opens a new empty dashboard tab with the name “Untitled”.
- 2 To place items onto the dashboard, add widgets by clicking the **Add Widget**  button located in the upper right corner. Select a widget and click-and-drag it onto the dashboard.
- 3 For each item (widget) placed, specify Widget Properties, as needed.



Note

By default, a scroll bar is not available in the Dashboard for external links. To include a scroll bar, set the “Show Scrollbar” property to “Yes” in the Widget Properties section of “External Links” under Dashboard Items.

- 4 Click **Save** to save the dashboard.



Note










Once saved, new dashboards become available in the **Dashboard Preferences** list of “Available Dashboard(s)”.

See [“Dashboards” on page 153](#) for information on how to display the new dashboard you just created or set the default display to a different dashboard.

Dashboard Buttons

The Dashboard Properties are described in the following table.

Table 5-1 Dashboard Properties Description

Button	Description
	Add widget Add a new Widget.
	Open page Open a page to edit the dashboard properties.
	Full screen view Shows the dashboard in a separate detached window.
	Open tab Open a new empty tab in the Dashboard Viewer where you can create a new dashboard.
	Refresh Now Refreshes the dashboard.
	Remove tab Dashboards are displayed in tabs in the Dashboard Viewer. Clicking this button will delete a tab from the Dashboard Viewer but will not delete the dashboard from its saved location. However, if you check the "Remove this dashboard from saved location" checkbox it will delete the Dashboard from the saved location too in addition to removing the dashboards tab from the Dashboard Viewer.
	Save Dashboard Saves an existing dashboard that has been modified.
	Save Dashboard As Saves the dashboard in your specified location.
	Subscribe Subscribes you to the Dashboard.

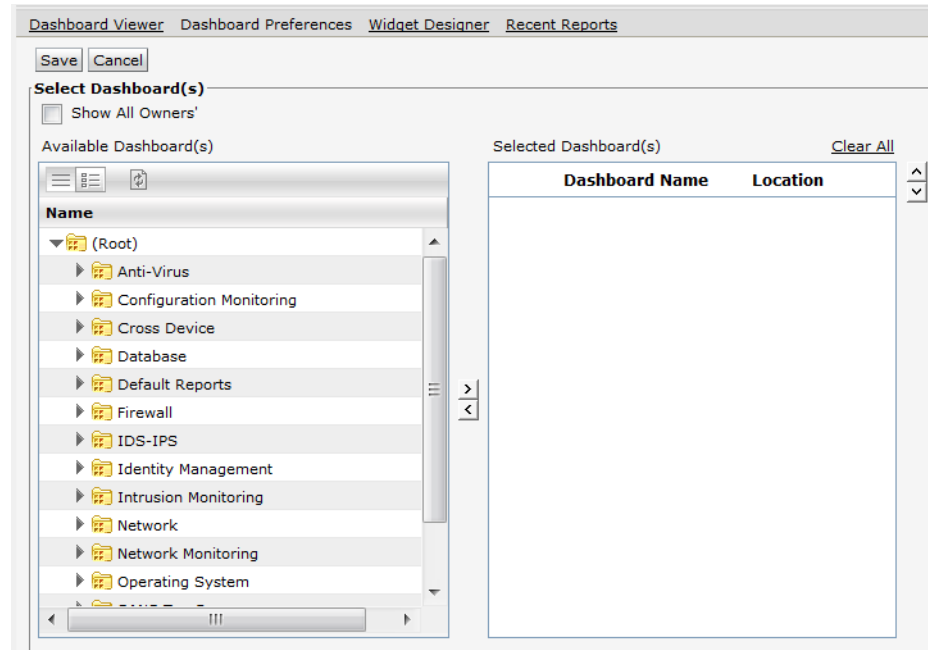
Viewing an Existing Dashboard in a Tab in the Dashboard Viewer


The set or subset of dashboards shown under Available Dashboard(s) is based on your user group status and the selection status of **Show All Owners'** checkbox. A user with Administrative rights is able to see more or all dashboards than a user with fewer privileges. If you limit the view to only your dashboards, the list will not include dashboards designed by other users.

- To access dashboards from all the users (designers), click (checkmark) the **Show All Owners'** checkbox.
- To view only your dashboards, click (uncheck) this checkbox.



To open multiple dashboards as tabs in the Dashboard Viewer:

- 1 Click the **Dashboard Preferences** link on top of the Dashboard Viewer page.
- 2 In the Available Dashboards box, navigate to the dashboard that you want to display in a tab.





- 3 Click the right-facing arrow  in between the two boxes. The dashboard name appears in the Selected Dashboard box.
- 4 Click the **Save** button.
- 5 Click the **Dashboard** link in the left panel to display the Dashboard Viewer. You should now see a tab displaying the dashboard you selected.


Removing an Existing Tab from the Dashboard Viewer

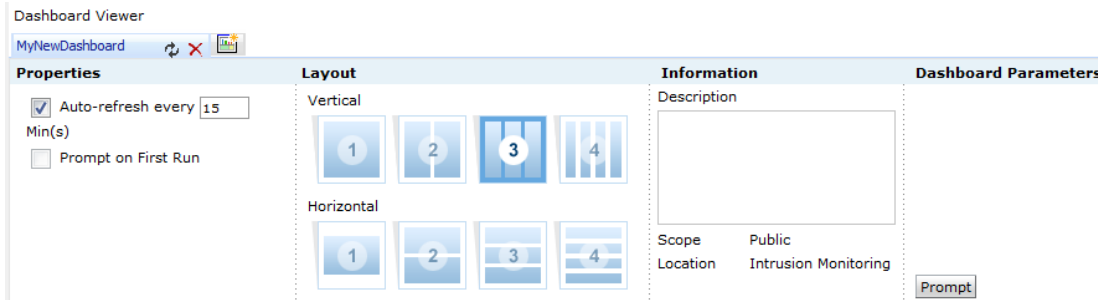
To remove an existing tab from the Dashboard Viewer without deleting the dashboard from its saved location, click the **Remove Tab**  button on the far right of the tab. To remove an existing tab from the Dashboard Viewer and delete the dashboard from its saved location, click the **Remove Tab**  button on the far right of the tab. In the Remove Tab dialog, check the **Remove this dashboard from saved location** checkbox and click **OK**.

Deleting a Dashboard

You can delete an existing dashboard from the Category Explorer or from the Dashboard Viewer. To delete a dashboard from the Category Explorer, click on **Dashboards** in the Repository column then click on the dashboard to select it and click the **Delete**  button. To delete the dashboard from the Dashboard Viewer, click the **Remove Tab**  button on the far right of the tab. In the Remove Tab dialog, check the **Remove this dashboard from saved location** checkbox and click **OK**.

Editing an Existing Dashboard

To modify an existing dashboard, click the **Edit Dashboard**  button in the Dashboard Viewer page. Its current configuration is displayed in the Properties, Layout and Information area, and you can modify then save settings as needed.



The screenshot shows the 'Dashboard Viewer' window with a tab for 'MyNewDashboard'. The interface is divided into four main sections:

- Properties:** Contains checkboxes for 'Auto-refresh every' (checked) and 'Prompt on First Run' (unchecked). A text input field for 'Min(s)' is set to '15'.
- Layout:** Shows two options: 'Vertical' and 'Horizontal'. Each option has four numbered icons (1, 2, 3, 4) representing different dashboard pane configurations. In the 'Vertical' section, icon 3 is selected.
- Information:** Includes a large text area for 'Description', and fields for 'Scope' (set to 'Public') and 'Location' (set to 'Intrusion Monitoring'). A 'Prompt' button is at the bottom right.
- Dashboard Parameters:** This section is currently empty.

To auto-refresh a dashboard at a regular interval, check the **Auto-refresh every** check-box and specify time in terms of minutes in Min(s). This will automatically refresh the dashboard after the set number of minutes. Check the **Prompt on First Run** checkbox to display the Input Parameter Form, which shows the values of the Dashboard parameters before reports are run from the dashboard for the first time after they have been displayed on the dashboard.

The Layout area allows you to select number of panes you want the dashboard to have.

The Information area displays Description, Scope and Location where the dashboard is saved.

Selecting a Default Dashboard View for the Reports Home Page

If you have multiple dashboards open in tabs in the Dashboard Viewer, you can set one of the dashboards to display as the default dashboard for the Reports home page.

To set a default dashboard:

- 1 Click the **Dashboard Preferences** link on top of the Dashboard Viewer page.
- 2 In the Selected Dashboards box, click the radio button against the dashboard that you would like to display as the default dashboard in the Reports home page. Click the up arrow (possibly multiple times) until that dashboard shows up at the top of the list.
- 3 Click the **Save** button.
- 4 Click the Dashboard link in the left pane and your selected dashboard will show as the default tab (the first tab).

The Dashboard Preferences page has the following fields:

Field	Description
Show All Owners	To display all dashboards made by all the users in the Available Dashboard(s) box, check the Show All Owners' checkbox.
Available Dashboards	This box shows a list of all dashboards that are available for display in the Dashboard Viewer.

Field	Description
Selected Dashboards	Move the dashboards you want to display in the Dashboard Viewer from the Available Dashboards list to the Selected Dashboards box. Dashboards appearing in this box will be displayed as tabs in the Dashboard Viewer.

Widgets

After a new dashboard is created, you will need to add one or more widgets to it to get it to display your reports or web links. Each dashboard item must be placed in its own widget for display on the dashboard. A widget is designed in the Widget Designer. A widget can be placed on multiple dashboards.

The Widget Designer

Clicking on the **Widget Designer** link located above the Dashboard Viewer opens the Widget Designer page. The Widget Designer allows you to create a new widget, save a widget, edit a widget, or delete a widget. You can place a report or a web link (an external link) into a widget. Each widget can contain only one object.

Creating a New Widget

To create a new widget, click the **Widget Designer** link located above the Dashboard Viewer. On the Widget Designer page, you can choose what you want to place in the widget, a report or a web link.

The screenshot shows the 'Widget Designer' tab in a web application. At the top, there are navigation links: 'Dashboard Viewer', 'Dashboard Preferences', 'Widget Designer' (active), and 'Recent Reports'. Below these are buttons: 'Add New', 'Save', 'Save As', 'Open', 'Delete', and 'Cancel'. The 'Widget Name' field is empty. Under 'Contents', the 'Report' radio button is selected, and the 'Web Link' radio button is unselected. A 'Pre-generated' button is visible. Below this, there are three sections: 'Report' with a dropdown menu showing '(Select Report)' and a 'Report' icon; 'By Job' with a dropdown menu showing '(Select Job)', a 'By Job' icon, and a checkbox labeled 'Look in User's All Jobs'; and 'In Category' with a dropdown menu showing '(Root)', an 'In Category' icon, and a checkbox labeled 'User's Working Folder'. At the bottom, the 'Widget Properties' section contains several settings: 'Report Format' set to 'HTML', 'Toolbar' set to 'Multipage', 'Instance Navigation' set to 'No', 'Auto Refresh' set to 'Yes', 'Refresh Interval' set to '15' with a unit of 'Min(s)', 'Width' set to '0', and 'Height' set to '0'.

You cannot run reports from a Dashboard view; you can only view results of previously saved, published reports. A refresh or auto-refresh on a dashboard simply updates the dashboard display with the most recently published results; it does not run the report. Therefore, reports on dashboards must be run, saved, and published in order for the report

data to be viewable on the Dashboard view. If a report on a dashboard has not been saved or published, its widget will display an error message on the Dashboard view the report data is not available to the dashboard.

Placing a Report

If you choose to place a Report, keep in mind that you can only add reports that have already been run and published.

To get last published instance of:

- A report: You need not make any selection in the Report field, By Job or In Category.
- A specific report: Navigate to the report in the Report field. You can leave the By Job and In Category fields blank.
- A report executed by a specific job: Navigate to the job in the By Job field. You can leave the Report field and the In Category field blank.
- A report deployed in a specific category and executed by a specific job: From the In Category field, navigate to a category and navigate to a job in the By Job field.
- Any of the reports from the jobs you own: You own the jobs that you created or were created on your behalf. Check the **Look in User's All Jobs** checkbox.
- Any of the reports deployed in your default category: Check the **User's Working Folder** checkbox.

Specify the following widget properties:

Label	Description
Widget Name	Enter a name for the new widget to be created
Report Format	Select the format in which you would like the report displayed.
Toolbar	Select whether you want a toolbar displayed and whether you want it displayed on all pages if this is a multi-page report.
Instance Navigation	<p>Sets whether to include a report navigation feature on the dashboard.</p> <ul style="list-style-type: none"> Click Yes to provide a drop-down menu that allows Dashboard users to select a saved report and view it. Click No if you do not want to provide this feature on the dashboard.
Auto Refresh	Set it to yes, if you want the report to refresh automatically after a certain interval. You must set the Refresh Interval parameter if you set Auto Refresh to Yes.
Refresh Interval	This is the time in minutes. Refresh will take place at the end of specified number of minutes. For example, if you want the report results to refresh every 15 minutes, set the Refresh Interval to 15.
Width	Select the width of the widget in pixels. You can select only whole numbers (no decimals allowed)
Height	Select the height of the widget in pixels. You can select only whole numbers (no decimals allowed)

Placing a Web Link

Click the Web Link radio button on the Widget Designer page to place a web link in the widget.

Specify the following properties:

Label	Description
URL	Specify the URL for the external link of the page that you want to display in the widget
Show Scrollbar	Select whether you want a scroll bar in the widget. By default the scrollbar is visible.
Auto Refresh	By default the web page will be automatically refreshed. Select No if you want to turn this feature off.
Refresh Interval	This is the time in minutes. Refresh will take place at the end of specified number of minutes. For example, if you want the web page to refresh every 15 minutes, set the Refresh Interval to 15.
Width	Select the width of the widget in pixels. You can select only whole numbers (no decimals allowed)
Height	Select the height of the widget in pixels. You can select only whole numbers (no decimals allowed)

Deleting a Widget


To delete a widget, open the widget in the Widget Designer and click the **Delete** button on top of the Widget Designer page. When a dashboard uses a widget that is deleted, an error message explaining that the widget has been deleted is displayed in the widget on the dashboard.

Editing an Existing Widget

Click the **Widget Designer** link to open it. Click the **Open** button on top of the Widget Designer page and select the widget that you want to edit. After editing it, click the **Save** button.

Placing Widgets in a Dashboard

Reports and Web Link (external link) objects are available to be placed on a dashboard. However, these objects must first be placed in a widget and then the widget can be added to the dashboard.

- 1 Click the **Add Widget**  button on the upper right corner of the empty dashboard page.
- 2 Navigate to the widget you want to place on the dashboard and click-and-drag it to the dashboard.
- 3 Continue to add more widgets by following the above two steps.

Moving an Existing Widget within a Dashboard

To move an existing widget on a dashboard, hover your mouse over the top boundary of the widget. The widget name bar will drop down. Click the widget name bar and drag it to move the widget to the desired location on the dashboard.

Using Dashboards Created in Pre-5.2 Logger

This section is applicable to existing pre-5.2 dashboards only. To create, edit or view new dashboards, follow the instructions in [“Dashboards” on page 153](#).

The ability to use and edit the pre-5.2 Dashboards is available for backward compatibility only.

Viewing a Classic Dashboard

To view the dashboard, click the **Classic Viewer** link on top of the Reports page. If no dashboard is configured and selected for display, the default Reports home page shows the **My Reports** page that lists the status of recently run or accessed reports.

To set a dashboard as the default Dashboard View:

Open the dashboard in the Classic Designer and flag the **Add to My Preferred List** checkbox.

The Classic Viewer page displays the contents of various items placed on the dashboard during the dashboard's design time. If the dashboard includes reports, reports will show current data from recently run reports.

Reports must be run and published first in order for the results to be accessible on a dashboard view. There are no options available to run reports from the Dashboard view. On a Dashboard view, you can view saved or published reports but not run them.

A refresh or auto-refresh on a dashboard simply updates the dashboard display with the most recently published results; it does not run the report. We suggest using the dashboard refresh interval in conjunction with scheduled reports to ensure that report results are always published and retained long enough to be available to dashboards.

Designing Classic Dashboards

Use the **Classic Designer** page to edit the dashboard and add items to it, and change the layout. Each dashboard can include multiple items (reports, use cases, and Web links).

What Items Can a Dashboard Include?


The following information is available for placement on a dashboard:

- **Reports**; any report can be included. The dashboard will show the latest published version of the report. Reports must be published in order for the report data to be accessible to users on the Dashboard View. If no published results are available for a report on a dashboard, the Dashboard View will display a message indicating this. When the report is published, a refresh of the Dashboard view will display the report.
- Common **Use Cases**, including a Report List, Saved Report List, Health Monitor, Recent Run Report List, Quick Job List, Schedule History, and Audit Log.
These are provided as dashboard elements so that users access a use case without leaving the Dashboard View page.
- **External Links**; that is, any URL(s) that you want on-screen as a part of a particular Dashboard View



Creating a New Classic Dashboard

Do not use the Classic Designer page to create new dashboards. New dashboards should be created using the **Dashboard Viewer** link. See [“Dashboards” on page 153](#) for more information.

Placing Items onto the Existing Dashboard

To place an item on the dashboard, click the **Classic Preferences** link. In the **Widgets** provided in the Layout area, click-and-drag an item from the **Dashboard Items** list on the left into an empty widget to the right. Alternatively, click  next to the dashboard item you want to place the item on an empty widget.

You can also click-and-drag an item onto a currently occupied widget if you want to replace an item in a widget with a different one.

To get a new widget, click  (Divide Widget Horizontally) or  (Divide Widget Vertically) on a widget to split it into two widgets. The original widget remains a new empty widget is placed on the dashboard layout.

For each item (widget) placed, specify Widget Properties, as needed.

By default, a scroll bar is not available in the Dashboard for external links. To include a scroll bar, set the “Show Scrollbar” property to “Yes” in the Widget Properties section of “External Links” under Dashboard Items.

Dashboard Properties



Dashboard Properties

Name:

☐ Public ☒ Private

Description:

Figure 5-1 Reports Dashboard Properties

The Dashboard Properties are described in the following table.



Table 5-2 Dashboard Properties Description

Property	Description
Name	Name of the dashboard.
Description	Descriptive information about this dashboard.

Creating Widgets

Each dashboard item must be placed in its own widget for display on the dashboard. Create a new widget using the **Widget Designer** link.

To get a new widget:

To get a new widget, click  (Divide Widget Horizontally) or  (Divide Widget Vertically) on a widget to split it into two widgets. The original widget remains a new empty widget is placed on the dashboard layout.

To remove a widget:

To remove a widget, click the Remove Widget button located on the top right corner on the widget you want to remove.

Placing Dashboard Items on the Layout

Click the **Classic Designer** link on the Reports page. Reports, use cases, and external link objects are available under “Dashboard Items” (to the left of the Layout area).




Figure 5-2 Dashboard Items

To place a dashboard item, click to expand the menu for the type of item you want, click-and-drag an item onto a widget in the Layout area, and specify widget properties as needed. (Widget properties vary depending on the type of item you place on the dashboard.)

The following sections provide more detail on placing each type of dashboard item and setting appropriate widget properties.

Placing a Report on a Dashboard

The following sections describe in detail how to place and configure reports on dashboards, including setting widget properties, report parameters, and dashboard parameters.

There are no options available to run reports from a Dashboard view; only to view results of previously saved, published reports. A refresh or auto-refresh on a dashboard simply updates the dashboard display with the most recently published results; it does not run the report.

Therefore, reports on dashboards must be run, saved, and published in order for the report data to be viewable on the Dashboard view. If a report on a dashboard has not been saved

or published, its widget will display an error message on the Dashboard view the report data is not available to the dashboard.

To place a report on a dashboard:

- 1 Under Dashboard Items, click **Reports** bar to expand the list of available reports.

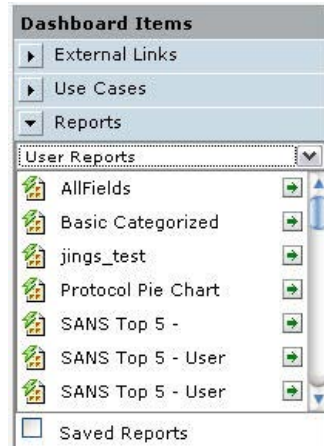



Figure 5-3 Reports under Dashboard Items

- 2 If available, select a Reports submenu such as **User Reports**, **Solution Reports**, and so forth.

Different reports are displayed depending on the submenu you select.

- 3 Optionally, check (select) **Saved Reports** checkbox to get a list of saved reports.
- 4 Select a category to view reports deployed in that category.
- 5 Click and drag the report to the widget in which you want to place the report. Alternatively, click  next to the dashboard item you want to place the item on an empty widget.

The report name is displayed in the widget in the Layout area.

- 6 Set Widget Properties for the report.

Widget Properties	
Report Name	SANS Top 5 -
Refresh Interval (in mins.)	15
Format	HTML
Auto Refresh	YES
Toolbar	MULTIPAGE
Instance Navigation	NO
Link Widgets	...
Description	

Figure 5-4 Widget Properties for Reports on a Dashboard

The following table describes Widget Properties settings for Reports dashboard items.




By default, a scroll bar is not available in the Dashboard for external links. To include a scroll bar, set the “Show Scrollbar” property to “Yes” in the Widget Properties section of “External Links” under Dashboard Items.

Table 5-3 Widget Properties for Reports on a Dashboard

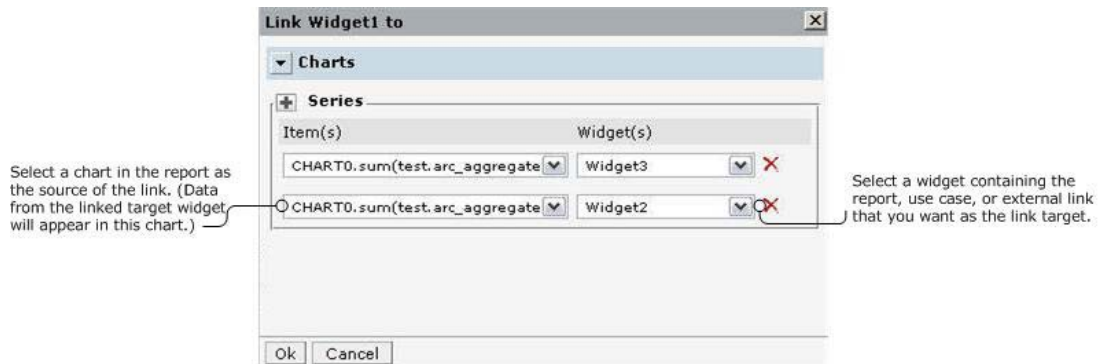
Property	Description
Report Name	The name of report that occupies this widget.
Refresh Interval (in minutes)	<p>This is the time in minutes. Refresh will take place at the end of specified number of minutes. For example, if you want the report results to refresh every 15 minutes, set the Refresh Interval to 15.</p> <p>Note: Reports must be run and published first in order for the results to be accessible on a dashboard view. A <i>refresh</i> or <i>auto-refresh</i> on a dashboard simply updates the dashboard display with the most recently published results; it does not run the report. We suggest using the dashboard refresh interval in conjunction with scheduled reports to ensure that report results are always published and retained long enough to be available to dashboards. (See also, “Scheduling Reports” on page 233.)</p>
Format	<p>Select the output format in which you want to view the report. Available options are:</p> <ul style="list-style-type: none"> • HTML • Acrobat PDF • Interactive
Auto Refresh	<p>Enables or disables auto-refresh option.</p> <ul style="list-style-type: none"> • Select Yes to refresh the reports as per Refresh Interval. • Select No to view the report generated when dashboard was loaded for the first time. <p>Note: Reports must be run and published first in order for the results to be accessible on a dashboard view. A <i>refresh</i> or <i>auto-refresh</i> on a dashboard simply updates the dashboard display with the most recently published results; it does not run the report. We suggest using the dashboard refresh interval in conjunction with scheduled reports to ensure that report results are always published and retained long enough to be available to dashboards. (See also, “Scheduling Reports” on page 233.)</p>
Toolbar	<p>Specifies Toolbar settings.</p> <ul style="list-style-type: none"> • Select Yes to always show toolbar. • Select No to never show the toolbar. • Select Multipage to show the toolbar only for multi-page reports. <p>The Multipage setting is applicable to HTML and Interactive output formats.</p>
Instance Navigation	<p>Sets whether to include a report navigation feature on the dashboard.</p> <ul style="list-style-type: none"> • Click Yes to provide a drop-down menu that allows Dashboard users to select a saved report and view it. • Click No if you do not want to provide this feature on the dashboard.

Table 5-3 Widget Properties for Reports on a Dashboard (Continued)



Property	Description
Link Widgets	Click  to bring up a Link Widget dialog in which you can specify a link from any of the charts in the report in this widget to another widget. See “Linking Widgets” on page 167 .
Description	Description of the widget.


Linking Widgets

You can link a widget that contains a report (although, not saved reports) to another widget. The widget that is the link target can contain a use case, a report, or external link.

**Figure 5-5** Linking Widgets

To link a chart in a report to data in another widget:

- 1 Select a widget in which you want to provide a link. (The widget that is the link “source” must contain a report with a chart on it).
- 2 Under Widget Properties for the selected widget, click  to bring up a Link Widgets dialog in which you can specify a link from any of the charts in the report to another widget. (The widget that is the target of the link can contain a report, use case or external link.)
- 3 In the Link Widget dialog, select an Item (chart series) from the Item(s) and select (link) it to an item in one of the other Widgets.
- 4 Click  (add button) next to “Series” to get another row to specify another set of link information in the same report with a different widget/series combination.

To remove a row, click  (delete button) next to the row you want to remove.


- 5 Click **OK** to save the settings and close the dialog.


Placing a Use Case on a Dashboard

The following sections describe in detail how to place and configure use cases on dashboards.

To place a use case on a dashboard:

- 1 Under Dashboard Items, click **Use Cases** bar to expand the list of available use cases.

**Figure 5-6** Use Cases under Dashboard Items

- 2 Click and drag a use case to the widget in which you want to place it. Alternatively, click  next to the dashboard item you want to place the item on an empty widget.

The use case name is displayed in the widget in the Layout area.

- 3 Set Widget Properties for the use case.

Widget Properties for Use Cases

Widget Properties	
Name	Health Monito
Refresh Interval (in mins.)	15
Auto Refresh	YES 
Show Scrollbar	NO 
Description	

Figure 5-7 Widget Properties for Use Cases on a Dashboard

The following table describes Widget Properties settings for Use Case dashboard items.

Table 5-4 Widget Properties for Use Cases on a Dashboard

Property	Description
Name	The name of use case that occupies this widget.
Refresh Interval (in minutes)	This is the time in minutes. The use case page is refreshed at the specified interval.
Auto Refresh	Enables or disables auto-refresh option. <ul style="list-style-type: none"> • Select Yes to refresh the use case as per Refresh Interval. • Select No to execute only once, when the dashboard is loaded.
Show Scroll Bar	Select Yes to get scroll bar if use case does not fit in widget width.
Description	Description of the widget.

Table 5-4 Widget Properties for Use Cases on a Dashboard (Continued)

Property	Description
Category	This option appears when Report List, Saved Report List or Quick Job List is placed on widget. Select the category to carry out respective task (get a list of reports in selected category, get a list of saved reports or quick job lists for selected report).
Report	This option appears when Saved Report List or Quick Job List is selected. Select the report for which saved report list or quick job list is to be viewed.


The use cases displayed in the list will depend on the permissions associated with your user group. Other properties are displayed based on the use case.


Placing an External Link on a Dashboard

The following sections describe in detail how to place and configure an external link on a dashboard.

To place a link on a dashboard:

- 1 Under Dashboard Items, click **External Links** bar to expand the list.

**Figure 5-8** External Link under Dashboard Items

- 2 Click and drag an External Link URL object to the widget in which you want to place it. Alternatively, click  next to the dashboard item you want to place the item on an empty widget.

The External Link URL object is displayed in the widget in the Layout area.

- 3 Set Widget Properties for the URL.

**Figure 5-9** Widget Properties for an External Link on a Dashboard

The following table describes Widget Properties settings for External Links dashboard items.

Table 5-5 Widget Properties for External Links on a Dashboard

Property	Description
Name	The name of use case that occupies this widget.
Refresh Interval (in minutes)	This is the time in minutes. The use case page is refreshed at the specified interval.
Auto Refresh	Enables or disables auto-refresh option. <ul style="list-style-type: none"> Select Yes to refresh the URL as per Refresh Interval. Select No to execute only once, when the dashboard is loaded.
Show Scroll Bar	Select Yes to get scroll bar if external link does not fit in widget width.
Description	Description of the widget.
URL	Specify the URL for this widget. If you want to add multiple Web pages to the dashboard, use a different widget for each URL.

Swapping Items on Widgets

You can swap items placed in widgets. To do this, click and drag the item to the widget where you want to place it.

Click and drag an item to a different widget to swap placement of the two items on the page.



Figure 5-10 Swapping Widgets on a Dashboard Design

In the above example, the Recent Run Reports List item is swapped to the position of the External Link URL, which is then swapped to with the Health Monitor item, which will end up at the top of the dashboard.

Setting Pre-5.0 Dashboard Preferences

In the **Classic Preferences** page, you can specify the dashboard to be made available for viewing.

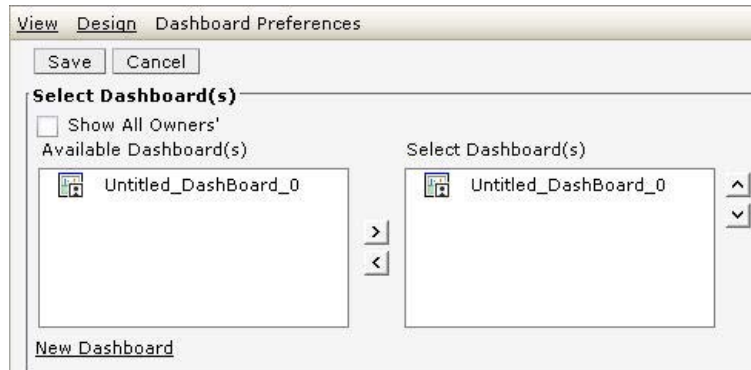


Figure 5-11 Dashboard Preferences

Working with Available Dashboards

The set or subset of dashboards shown under Available Dashboard(s) is based on your user group status. For example, it is likely that a user with Administrative status will be able to see more or all dashboards than a user with fewer privileges.

Selecting a Dashboard View


Once you have created one or more dashboards, you can select one of them as the default display for the Dashboard **View** page, which also serves as the Reports home page. Only one dashboard at a time can be displayed as the default dashboard view.




You must have at least one dashboard in order to set a preference for the Dashboard View.

You can also set a dashboard as the “Selected Dashboard” (default dashboard view) in the Dashboard Designer by enabling the **Add to my preferred list**, as described in [“Viewing a Classic Dashboard” on page 162](#).

To select a default Dashboard View for the Reports home page:

- 1 Navigate to **Dashboard > Preferences**.
- 2 Select a dashboard from the Available Dashboard(s) list and click the right arrow button  to move it into the Select Dashboard(s) list for display. Only one dashboard can occupy the “Selected Dashboard(s)” list at any one time.
- 3 Click **Save** to save your preferences and display the selected dashboard.

To remove or change the currently displayed dashboard:

- 1 Return to the **Classic Preferences** page.
- 2 Move the currently selected dashboard out of the Select Dashboard(s) list by selecting it and clicking the left arrow button .
- 3 Choose a different one to display if so desired (or none).
- 4 Click **Save** to save your preferences.

Running, Viewing, and Publishing Reports

Reports are deployed (made available) under their respective categories. (See [“Report Explorer” on page 141](#))

You can run, view, and publish reports you create, as well as reports in categories for which the administrator has given you user access rights. You can run up to five reports concurrently on a Logger.



There are no options available to run reports from a Dashboard viewer. On a Dashboard viewer, you can view saved or published reports but not run them.

You can run, view, customize or publish reports in the following ways:

- Clicking the Report Explorer link, clicking on a category, selecting the report and clicking on a desired button on the button bar. OR
- Double-clicking on a category in any explorer, then selecting the report by clicking the radio button next to it, and then clicking on a desired button on top of the page.

Best Practices

Logger is designed to process events while running a report, but event processing has priority. Running a complex report while the event processing system is under load will result in report timeout rather than dropped events.

HP recommends using the Scheduled Report feature so that reports are run during periods of light load. If an ad hoc report must be run, run it when the system is not under load.

For information on working with scheduled reports, see [“Scheduling Reports” on page 233](#).

If you are running a distributed report, also see the best practices discussed in [“Selecting Device Groups, Storage Groups, Devices, or Peers” on page 177](#).

Finding Reports

You can find reports on the following pages within the **Reports** tab:

- Category Explorer
- Report Explorer
- Favorites Explorer (if you have marked the Report as a Favorite)
- Scheduled Reports page (If the report you are looking for is a scheduled report and it has been run and published)
- Recent Reports link on top of any Reports page.

Viewing Recently Run Reports

To view the most recently run reports click the Recent Reports link on top of the Reports page.

Dashboard Viewer | Dashboard Preferences | Widget Designer | Recent Reports | Classic Viewer | Classic Designer | Classic Preferences

Recent Reports

Sr.No.	Report Name	Category Name	Time Stamp
1	<input type="radio"/> VPN Connection Attempts	Network Monitoring	11/03/2011 15:51:23
2	<input type="radio"/> IDM-Connection Counts by User	Identity Management	11/03/2011 15:05:59
3	<input type="radio"/> Most Common Events	Intrusion Monitoring	11/03/2011 14:59:55
4	<input type="radio"/> VPN Connection Attempts	Network Monitoring	11/03/2011 14:59:26
5	<input type="radio"/> AV-Update Summary	Anti-Virus	11/03/2011 14:59:01
6	<input type="radio"/> XD-BW-Top Bandwidth Hosts	Cross Device	11/03/2011 14:58:32

Report Execution Status

Filters: Category Name [(All)] | Report Name [(All)] | Execution Type [(All)] | Status [(All)] | User [ArcSight/admin]

Select Report(s): (Root)

Execution Type: (All) Status: (All)

Select Owner: ArcSight admin

Date From: 11/03/2011 To: 11/03/2011

Refresh

Sr.No.	Report Name	Action	Execution Type	Completion Status	Completion Date
1	<input type="radio"/> Network Monitoring/VPN Connection Attempts	VIEW	Run	Success	11/03/2011 15:51:23
2	<input type="radio"/> Identity Management/IDM-Connection Counts by User	VIEW	Run in Background	Success	11/03/2011 15:05:59

The Recent Reports box shows you the reports that were most recently run. You can click on the radio button next to a report to select it. After you select a report the **Run** and the **Re-Run** buttons appear on the top left corner. You can run the selected report using the same filter options as the original run by clicking on the Run button or you can run the selected report using different field values by clicking on the Re-Run button. See [“Run Report Parameters” on page 178](#) for details.

Task Options on Available Reports

Your access to various reports and report options (view, publish, edit, etc.) depends the access rights associated with your user and Logger Report Group affiliation. For example; depending on your access rights, you may have privileges to view some reports but not others, to view but not schedule or publish a report, or to view and schedule but not edit a report.

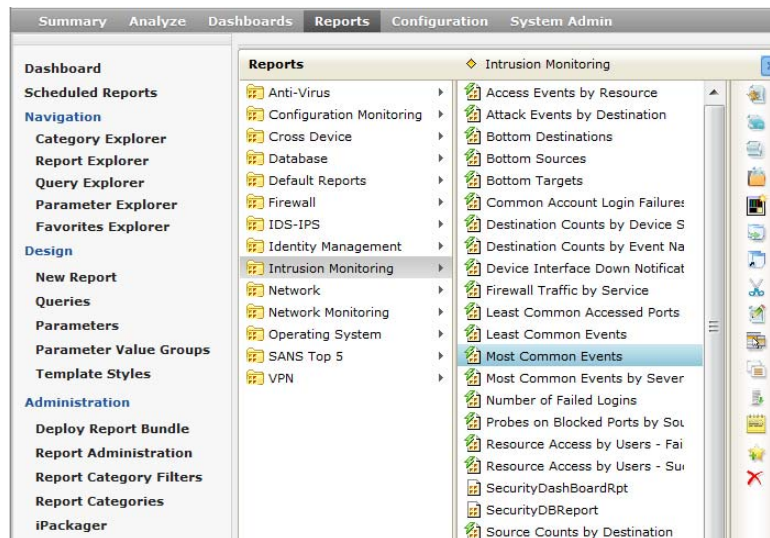
Access rights to report options are configured and managed with the **User/Groups** option on the **System Admin** page. For more information on setting permissions and on Logger Report Group management, see [“Setting Access Rights on Reports” on page 198](#).

See [“List of Buttons in the Explorers” on page 144](#) for buttons that represent various actions that can be done on Reports.

The following sections describe details of running and viewing reports, setting report parameters on a “Quick Run”, “Run in Background”, or “Run” of a report, and the various options for working with report output.

Running and Viewing Reports

To get started running and viewing reports, choose a report category in the Reports Explorer, and then choose a report within the category.



About the Pagination of Reports

The default view option for the report results is Multipage. If you uncheck the Multipage option in View Options link, your report will be formatted as a single page. **HP strongly recommends using the Multipage option for all reports.** Specifically, if a report will result in more than 826 records, using the Multipage option will ensure that the generated report is not blank. By default, the reports generated in the PDF format are set to use the Multipage option. However, if your PDF format report is blank, ensure that the Multipage option is still checked for that report.

If a report contains more columns than can be displayed horizontally across a screen using the default width specified in the report query, the report is paginated horizontally such that additional columns are displayed on the following pages. For example, if a report contains 45 columns and only 5 can be displayed on each screen, the report would be paginated such that Page 1 displays columns 1 through 5, Page 2 displays columns 6 through 10, Page 3 displays columns 11 through 15, and Page 9 displays columns 40 through 45. Consequently, if the report contained more rows than can be displayed vertically in one screen, the second screen of rows would be displayed starting at Page 10.

Currently, Logger limits the number of pages for horizontal pagination to 10. Consequently, if a report requires more than 10 pages to display all columns, complete report results may not be displayed. To view all columns of such reports, manually set the width of each column such that all columns fit in 10 or less pages in the report query (in Query Object Editor that can be opened by clicking the Queries link in the left pane).

To run and view a report:

Reports can be run from any of the following explorers:

- **Category Explorer**

Click a category in the Repository column, click Reports, select a report in the Reports column and click either Quick run with default options, Run in Background, or Run Report buttons. Set the parameters and click Run Now, Run in Background, or Run button depending upon what type of run you chose.

- **Report Explorer**

Click a category in the Reports column, select a report in the next column, and click either Quick run with default options, Run in Background, or Run Report buttons. Set the parameters and click Run Now, Run in Background, or Run button depending upon what type of run you chose.

- **Favorites Explorer (if you have marked the Report as a Favorite)**

Click on a report in the Favorites column and click either Quick run with default options, Run in Background, or Run Report buttons. Set the parameters and click Run Now, Run in Background, or Run button depending upon what type of run you chose.



Even if you selected Run Report initially, you can run a report in the background after setting the Run Report parameters.

The report output is displayed in the specified format (HTML, PDF, or other).

At this point, the results of this report generation is available as a file for viewing only by you. If you close the file without saving or publishing it, the results are no longer available.

If you want to make the results of this run available for others, publish it. To do this, leave the file open, click the Publish Report button available in the button bar located on top of the report, and follow the steps in [“Publishing Reports” on page 180](#).

For information about other delivery options available to you at this point, see [“Report Delivery Options” on page 181](#).

Quick Run with Default Options / Run In Background Report Parameters

When you click the **Quick Run with default options**  button or the **Run in**


Background  button for a report, the report will run with the data filters specified in the deployed report. You still get options to select additional filters on timeframe and constraints—Device Groups, Storage Groups, Devices, and Peers—over which the report runs. Starting with Logger 5.2, you can run a distributed report—a report that also includes matching events from the specified peers of a Logger. You select the peers on which the report should run in the Peers list, as shown in the following figure. If no peers are configured, the Peers list contains only the localhost IP address (127.0.0.1); however, if peers are configured, their IP addresses are listed.

Figure 5-12 “Quick Run with default options” / “Run in Background” Report Parameters

The following table describes Quick Run with default options with default options / Run in Background report parameters.

Table 5-6 “Quick Run with default options” / “Run in Background” Report Parameters

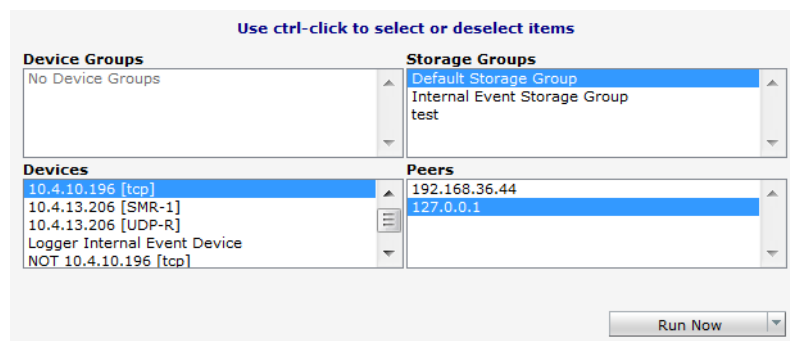
Option	Description
Start	<p>Specify the starting point for the data gathering from the events database.</p> <p>By default, the start time is specified with a dynamic data expression (\$Now - 2h).</p> <p>You can modify the dynamic expression to specify a different dynamic start time, or disable Dynamic and use the calendar options to specify a fixed start time.</p>
End	<p>Specify the ending point for the data gathering that is some time after the starting point.</p> <p>Keep in mind that large time spans can mean large amounts of data, which can affect system performance.</p> <p>By default, the end time is specified with a dynamic data expression (\$Now).</p> <p>You can modify the dynamic expression to specify a different dynamic end time, or disable Dynamic and use the calendar options to specify a fixed end time.</p>
Scan Limit	<p>Specify the number of events to scan.</p> <p>When you specify a scan limit, the number of events scanned for <i>manually</i> run reports is restricted to the specified limit. Doing so results in faster report generation and is beneficial in situations when you only want to process the latest N number of events in the specified time range instead of all the events stored in Logger.</p> <p>The scan limit is 100,000 by default. If you set the scan limit to 0 (zero), all events are scanned.</p> <p>This setting does not apply to the scheduled reports.</p>
Device Groups	<p>Select the device group(s) on which to run the report query, if any. (See “Selecting Device Groups, Storage Groups, Devices, or Peers” on page 177.)</p>
Storage Groups	<p>Select the storage group(s) on which to run the report query. (See “Selecting Device Groups, Storage Groups, Devices, or Peers” on page 177.)</p>

Table 5-6 “Quick Run with default options” / “Run in Background” Report Parameters

Option	Description
Devices	Select the device(s) on which to run the report query. (See “Selecting Device Groups, Storage Groups, Devices, or Peers” on page 177.)
Peers	Select the peer Loggers on which to run the report query. If no peers are configured on the Logger, this option only lists the localhost IP address (127.0.0.1); however, if peers are configured, their IP addresses are listed. (See “Selecting Device Groups, Storage Groups, Devices, or Peers” on page 177.)

Selecting Device Groups, Storage Groups, Devices, or Peers

The following figure shows how to select or de-select items on Device Groups, Storage Groups, Devices, or Peers as a part of setting Report “Quick Run with default options” and “Run in Background” parameters.

**Figure 5-13** Selection Model for “Quick Run with default options” or “Run in Background” Report Scope of Storage and Devices


- Items with a blue highlight are selected and will be included in the report query when the report is run.
- Items that are not highlighted are de-selected and will not be included in the report query.
- To select an item, click on it. To select multiple items in a list, use Ctrl-Click.
- To de-select a currently selected item, use Ctrl-Click.
- If none of the storage groups, device groups, or devices are selected, all items are included in the report query. However, peers must be explicitly selected to run a report query on them; if none of the peers are selected, the query will only run on the local Logger.
- The selected items in the Device Groups, the Devices lists, and Peers are ORed in the report query, and these items are ANDed with the other selected items such as Storage Groups.
- Follow these guidelines when you run a distributed report (a report on peer Loggers):
 - ◆ You can run a distributed report in the Quick Run, Run in Background, Run, or Scheduled Reports mode.
 - ◆ All Loggers on which you are running the distributed report must be running Logger 5.2 or later. Peer Loggers can be of mixed form factors, that is, software and appliance based.
 - ◆ If peer Loggers do not have identical storage or device group names, the report query skips searching for events for those groups on those peers.

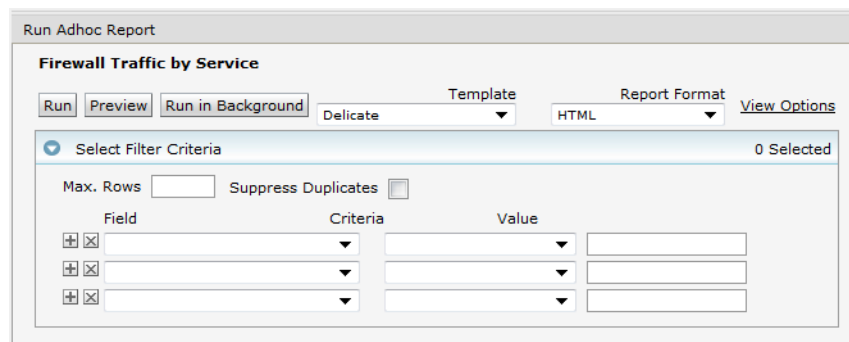
- ◆ If you added custom schema fields to your Logger schema, those fields must exist on all peers. Otherwise, a query containing those fields will not run (when run across peers) and return an error. See [“Adding or Importing Schema Fields” on page 345](#).
- ◆ A user needs to belong to these user groups with the listed permissions set to run distributed reports:
 - Logger Search Group with “Search for events on remote peers” user right set (checked).
 - Logger Rights Group with the “View registered peers” user rights set (checked).
 - Logger Reports Group with “View, run, and schedule reports” rights set (checked) for specific reports or the global permission set to run all reports.

For more information on setting permissions, see [“Setting Access Rights on Reports” on page 198](#)

- ◆ If a peer is unavailable when a distributed report is run, an error message is displayed and the report is aborted. Similarly, if a peer becomes unavailable while a distributed report is running, the report will continue to run and displayed; however, the server log will contain exceptions indicating the cause.
- ◆ Use the following best practices for optimal performance when running a distributed report:
 - Avoid running a distributed report on a Wide Area Network (WAN) link.
 - If you are running the report on a very large data set and the performance of the report is not optimal, reduce the size of the dataset.
 - Ensure that all fields in the report query are indexed on all peer Loggers. The report query will run slower on the Logger on which the fields are not indexed.

Run Report Parameters

When you click the **Run Report**  button for a report, you get additional options (beyond what you get for a Quick Run or for Run in Background) to choose a file format, specify pagination, and to modify the data filter criteria for only this run of the report.



Run Adhoc Report

Firewall Traffic by Service

Run Preview Run in Background Template: Delicate Report Format: HTML View Options

Select Filter Criteria 0 Selected

Max. Rows: Suppress Duplicates: ☐

Field	Criteria	Value
<input type="checkbox"/> <input type="checkbox"/>	<input type="text"/>	<input type="text"/>
<input type="checkbox"/> <input type="checkbox"/>	<input type="text"/>	<input type="text"/>
<input type="checkbox"/> <input type="checkbox"/>	<input type="text"/>	<input type="text"/>

If you run the report without specifying any override run-time parameters here, the report is generated with the defaults specified at design time for this report. You can run a report in the background after specifying the Run Report parameters.

The following table describes “Run” report parameters.

Table 5-7 Run Report Parameters

Option	Description
Report Format	<p>Specify a file type or “format” option of the output, and toggle on or off the Multipage option to generate a report as a multi-page or a single-page document. By default, Multipage is checked.</p> <p>Note: HP strongly recommends using the Multipage option for all reports. Specifically, if a report will result in more than 826 records, using the Multipage option will ensure that the generated report is not blank. By default, the reports generated in the PDF format are set to use the Multipage option. However, if your PDF format report is blank, ensure that the Multipage option is still checked for that report.</p> <p>For descriptions of report format see “Report File Formats” on page 179</p>
Select Filter Criteria	<p>Provides options to define filters, or modify default filters if any are already built in to the report.</p> <p>The filter expression is applied when the report runs, narrowing the focus of the report to the specified criteria.</p> <p>For example, you could set the filter criteria on a report on Top Password Changes to report only on password changes related to specified usernames or involving specified IP addresses.</p> <p>For details on how to create these filters (with Field, Criteria, and Value fields), see “Select Filter Criteria” on page 189 in “Designing New Reports” on page 187.</p> <p>Note: Filter criteria defined at report run time applies only to this run of the report. Filters set in this way are not saved nor made available to other users. You can also set built-in, default filter criteria as a part of designing a report.</p>
Template	<p>Select the template to apply to this report. The templates drop-down menu shows supplied templates, and any custom templates you may have added. These templates define the look and feel, arrangement, orientation, and so on, of the report output. To include the start time, end time, scan limit, device group, storage group, and devices information (used to run a report) in a report, choose the “BlankWithHeader” template.</p> <p>See “Applying Report Template Styles” on page 232 for more information on working with templates.</p>
View Options	<p>The default view option for the report results is Multipage. If you uncheck the Multipage option, your report will be formatted as a single page. HP strongly recommends using the Multipage option for all reports. You also have the option to download the report results as a zipped file. To do so, check the Download Zipped option.</p>

When you click **Run** on this first “Parameters” dialog, you then get the same dialog you get for a Quick Run with default option (or Run in Background) report where you can specify filters on timeframe and storage groups on which to run the report. (See [“Quick Run with Default Options / Run In Background Report Parameters” on page 175](#) for details on this “Additional Filters” dialog. Clicking **Run Now** on this second dialog runs the report.

Report File Formats

Report file formats include:

- HTML (Web page format)
- PDF (Acrobat PDF)
- MS Excel

- Comma Separated (Delimiter separated file. The delimiter is usually a comma.)
- MS Word
- Interactive (iHTML)
- XML

For most formats, you can select Multipage option by clicking on the **View Options** link. **HP strongly recommends using this option for all reports.** (If this option is checked, the report results will be formatted for a multi-page report.)

The report formats made available to you depend on access rights associated with your user account. (See [“Setting Access Rights on Reports” on page 198](#) for more information.)


Some report formats require that the workstation have respective Viewers. For example, PDF format needs Adobe Reader.

Publishing Reports

If you publish a report after you run it ([“Running and Viewing Reports” on page 174](#)), the output results for that run of the report are saved for subsequent use.

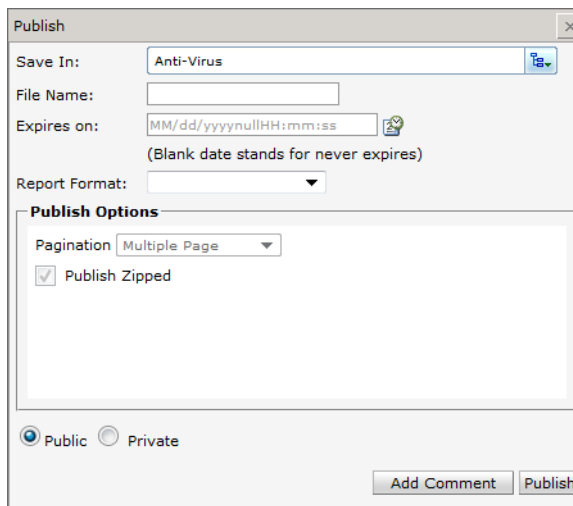
You configure *scheduled reports* to publish after each scheduled run. The publish options for scheduled reports are the same as for *on-demand reports* described here. For more about scheduled reports, see [“Scheduling Reports” on page 233](#) and [“Scheduling Reports” on page 233](#) and [“Add Report Job Settings” on page 236](#).

To publish a report:

- 1 In a generated report output file you get from running a report, click the **Publish Report**  button at the top of the page.

This brings up a Publish Report dialog in which to specify a file name for the report output, an expiration time if needed, and public or private status.

- 2 Specify the details with which to publish the report.



Publish

Save In: Anti-Virus

File Name:

Expires on: MM/dd/yyyynullHH:mm:ss
(Blank date stands for never expires)

Report Format:

Publish Options

Pagination: Multiple Page

☒ Publish Zipped

☒ Public ☐ Private

Add Comment Publish

The following table describes the publish report options.

Table 5-8 Publish Report Settings

Option	Description
Save In	Category under which to save the report. If you specify a category in the preferences, you can navigate to it. If you had not specified a category, the published report will be saved in the category in which the report resides. Note: You cannot save a report in the root category. Save it in one of the existing categories. Alternatively, save it in a new category.
File Name	Name for this report on the published reports list.
Expires on	Date and time after which the report output discarded (and, therefore, unavailable for viewing). If you do not want the report results to expire (keep always available), then leave this field blank (that is; do not set an "Expires on" date/time).
Report Format	Format of the report
Public or Private	Setting this as Public makes this report available to everyone. Setting this as Private makes this report available to you only.






3 Click **Publish**.













For information on how to view a published report, see ["Viewing the Output of a Published Report" on page 184](#).

Report Delivery Options

When you run a report from the Report Explorer (as described in ["Running and Viewing Reports" on page 174](#)), many options are available for delivering the generated output.

The most common next step is to publish the resulting report (described in ["Publishing Reports" on page 180](#)), but you can also save the report output to a file, e-mail it to other users, refresh the results, change the output format, and so forth. The following buttons are available from a generated Report representing various actions:

Button	Description
	Add Comment
	Refresh Comments
	Show Comments
	TOC
	MS Excel

Button	Description
	PDF
	CSV
	MS Word
	Export
	Email Report
	Upload Report
	First Page
	Previous Page
	Next Page
	Last Page
	Publish Report
	Refresh

Refreshing a Report


To re-run the report and get an updated result set, click  (Refresh).

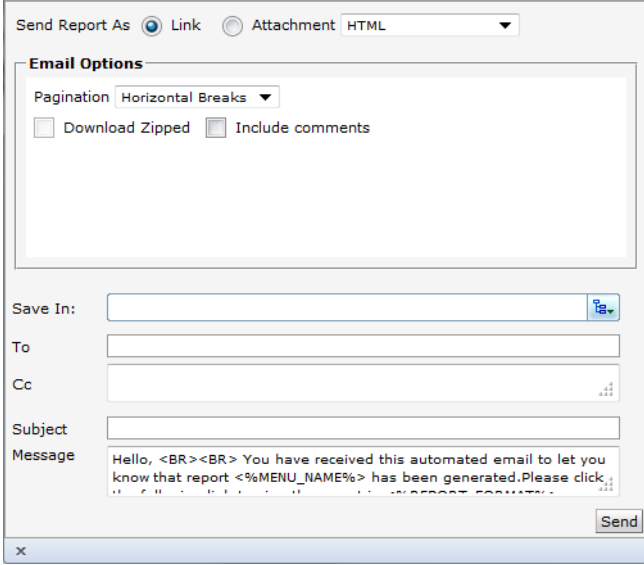
E-mailing a Report

You can send a report via e-mail as either a Web link or an attachment. You can also configure these e-mail options on *scheduled reports*, as described in [“Scheduling Reports” on page 233](#) and [“Add Report Job Settings” on page 236](#).

Before you can email a report, you must first set up SMTP for reports. You can do this by using the **SMTP Server** setting on the **Report Configuration** page as described in [“Report Configuration” on page 238](#).

To e-mail a report:

- 1 In the Report Explorer, after you run a report, click the  (Email report) button on top of the page.
- 2 Specify the following information about the e-mail.


Figure 5-14 E-mail Report Settings

The following table describes the e-mail report options.

Table 5-9 E-mail Report Settings





Option	Description
Send Report As	Choose one of these: <ul style="list-style-type: none"> To provide a link to the report in the body of the e-mail, select Link. To send the report as an attachment to the e-mail, click Attachment, and select a format for the attachment file.
Email Options	Select the following options before attaching a report file to an email: <p>Pagination: Select whether you would like the report to appear in a single page, multiple pages and whether you want horizontal breaks.</p> <p>Download Zipped: Zip the file before attaching it with the mail</p> <p>Include comments: Whether you want any comments that were added to the report included</p>
Save In	You have the option to save the report in a location that you can specify here. <p>Note: You cannot save a report in the root category. Save it in one of the existing categories. Alternatively, save it in a new category.</p>
To and CC	Specify e-mail addresses to which to send the report.
Subject	Provide e-mail Subject header.
Message	For the body of the e-mail, you can use the default message provided, modify it, or enter your own message.

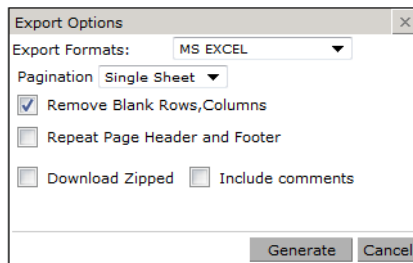
- 3 Click **Send** to send the report.

Exporting and Saving a Report

You can export a report to a file format of your choice and save it.

To export and save a report:


- 1 Click the **Export** button or click one of the file format buttons on the published report top-level menu bar (   ).
- 2 In the Export Options dialog, specify the Export Format and associated settings you want in the Export Options dialog.



Depending on the Export Format you choose, other settings are displayed as appropriate. Configure the export, and click **Generate**.

- 3 When the report is displayed, you have the option to save it as a file locally or elsewhere just as you would any other file.

Viewing the Output of a Published Report

- 1 Navigate to the report for which you want to view output results. (See [“Finding Reports” on page 172](#) if you need help locating a report.)
- 2 Click the **List Published Output button**  (Navigate to list of published outputs for this report) next to the report you are interested in.










Filters Published Name Includes [] Report Name [Intrusion Monitoring]/[Most Common Events]					
Intrusion Monitoring					
Sr.No.	File Name	Generated By	Generated Time	Expiry Time	
1		 ArcSight/admin RECENT	10/17/2011 16:22:53	Never	
2		 ArcSight/admin	10/12/2011 13:53:17	Never	

Figure 5-15 List of Published Report Outputs for a Selected Report


- 3 Click the radio button next to any of the reports listed to select options including the following:
 - View report outputs in various formats (HTML, PDF, CSV, Excel, or Word     )
 - Delete the selected instance of the generated report

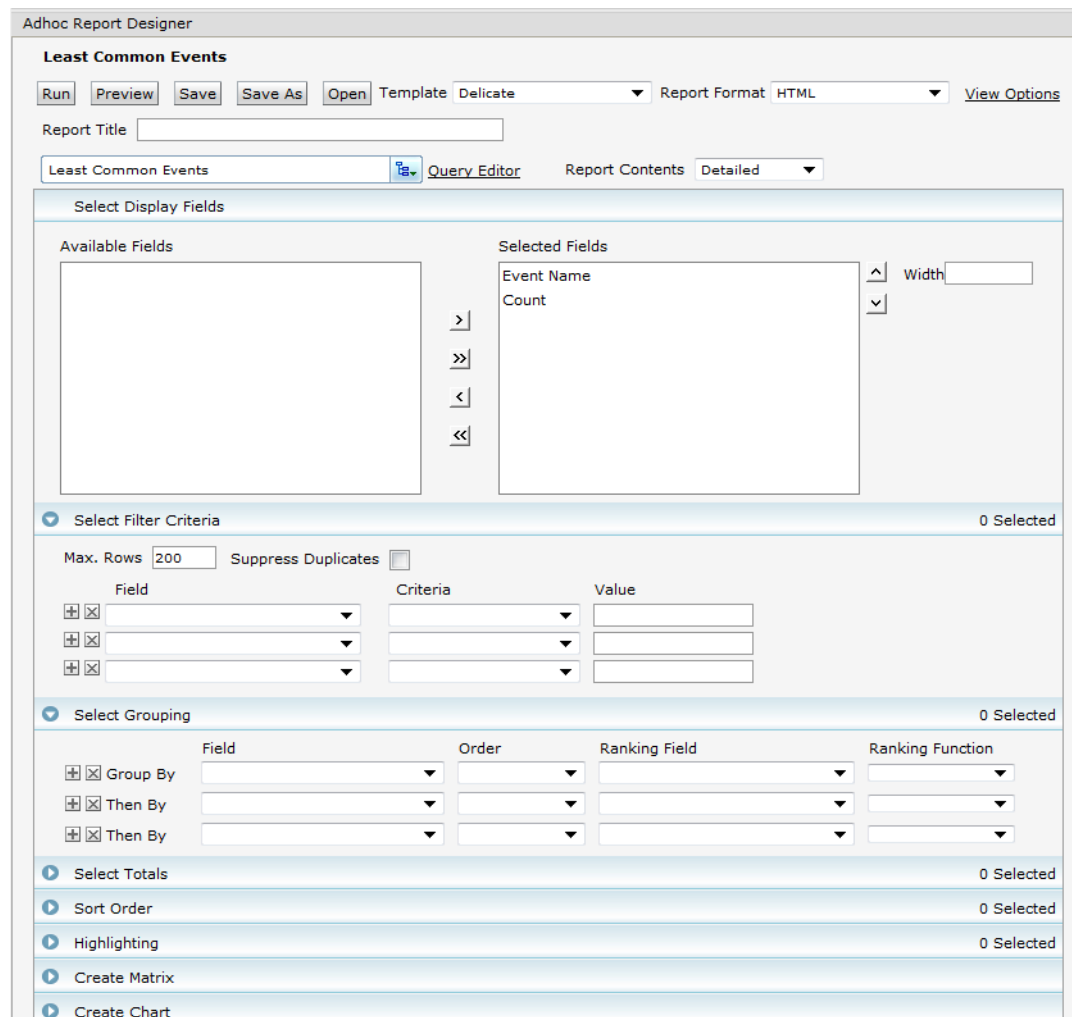
Designing Reports

You can use the Logger Report Designer to design simple columnar reports as well as mixed reports with embedded charts and matrices. For columnar reports, the Report Designer provides options for setting up filters, grouping, totals, and sort order to create a full-featured report.

Opening the Report Designer

To open the Report Designer to create a new report from scratch, click Design | **New Report** on the Reports left menu bar.

To open the Report Designer to edit an existing report, click the **Customize Report**  button for a report in a reports list in the Report Explorer. (See [“Explorers” on page 139](#) and [“Task Options on Available Reports” on page 173](#) for more information on available reports and how to get to their task option buttons, respectively.



The screenshot shows the 'Adhoc Report Designer' window. At the top, there's a 'Least Common Events' section with buttons for 'Run', 'Preview', 'Save', 'Save As', 'Open', and a 'Template' dropdown set to 'Delicate'. The 'Report Format' is set to 'HTML', and there's a 'View Options' link. Below this is a 'Report Title' text box. A 'Query Editor' button is next to a dropdown showing 'Least Common Events'. The 'Report Contents' dropdown is set to 'Detailed'.

The main area is divided into two panes: 'Available Fields' (empty) and 'Selected Fields' (containing 'Event Name' and 'Count'). Between them are navigation arrows (>, >>, <, <<). To the right of the 'Selected Fields' pane is a 'Width' input box.

Below the field panes is the 'Select Filter Criteria' section, which is currently empty (0 Selected). It includes a 'Max. Rows' input set to '200' and a 'Suppress Duplicates' checkbox. Below this is a table for filter criteria with columns for 'Field', 'Criteria', and 'Value'.

Next is the 'Select Grouping' section (0 Selected), which includes a table for grouping with columns for 'Field', 'Order', 'Ranking Field', and 'Ranking Function'. It has rows for 'Group By', 'Then By', and another 'Then By'.

At the bottom are several expandable sections: 'Select Totals' (0 Selected), 'Sort Order' (0 Selected), 'Highlighting' (0 Selected), 'Create Matrix', and 'Create Chart'.

Figure 5-16 Report Designer

Creating New Reports

If you are new to the Logger Report Designer, we recommend starting with an existing report as a basis for a new one, as described in [“Quick Start: Base a New Report on an Existing One” on page 186](#).

If you are starting a new report from scratch or for more details on each of the settings in the Report Designer, see [“Designing New Reports” on page 187](#).

Quick Start: Base a New Report on an Existing One

Since Logger ships with a variety of useful, pre-built reports for common security scenarios, you can leverage these not only to run as-is but also as templates for building new reports. If you are just getting started with the Report Designer, a good way to get up-to-speed fast is to start with an existing report that has some of the features you want in your new report, save the original report under a new name, and then modify it.

**Caution**

Modifications to reports and other ArcSight-defined content may be overwritten without warning when the content is upgraded. Do not modify ArcSight-defined content directly.

Make modifications to a copy of any ArcSight-defined content as a general practice, and subsequent upgrades will not affect the modifications.

To create a new report based on an existing report:

- 1 Navigate to the report you want to use as a starting point. (Click on a category in the Report Explorer to see a list of available reports in that category, and then click on the report.)

- 2 Click the **Customize Report**  button for a report in the Report Explorer toolbar.

This opens the Report Designer with the report's properties filled in.

**Note**

Some reports, such as the ones obtained from ArcSight or other custom developer sources might not be editable. For such reports, the icon is gray.

- 3 Click the **Save As** button.

This brings up the Save Report Layout As dialog for the selected report (and shows all reports stored in the same category as the one you selected).

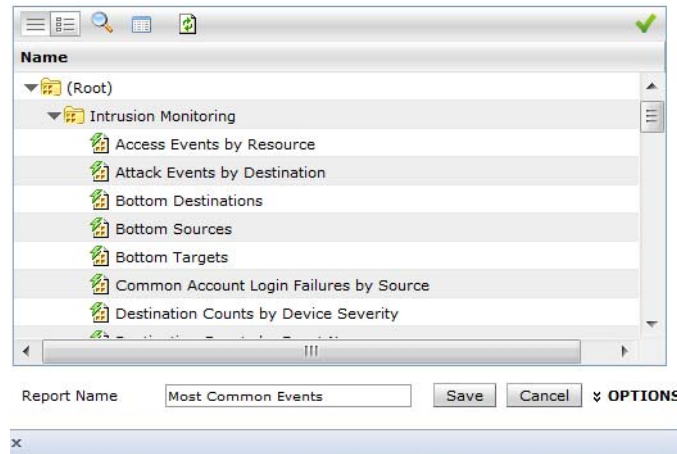


Figure 5-17 Save Report Layout As dialog for an Existing Report

- 4 Provide a Report Name for your new report.
- 5 Click **OPTIONS**.
Select **Public** (if you want everyone to have access to the report) or **Private** (to make the report available only to you), select **System Generated** if you do not want to enter a custom ID for the report and add a **Description**, if needed.
- 6 Click the **Save** button to save the report.
- 7 Click **OK** on the confirm dialog telling you that the report was saved successfully. Your new report is now shown under the category in which you saved it.
- 8 Select the report you just saved and click the **Customize Report** button to start modifying the new report to suit a specific scenario. (See the next section, [“Designing New Reports” on page 187.](#))

Designing New Reports

To access the Report Designer to create a new report from scratch, click Design | **New Report** link in the left panel.

This brings up the Report Designer with a blank template.

The following sections explain how to use the Report Designer.

Report Save, Run, and Template Options

- Click **Run** to test the current version of the report.
- Click **Preview** to preview the report before saving it.
- Click **Save** to save the report.
- Click **Save As** to save it under a different name.
- Click **Open** to open another report in the Report Designer.

General Report Settings

Set your preferences for pagination, layout and report output format as described below.

Table 5-10 General Report Design Settings

Option	Description
Template	<p>Select the template to apply to this report. The templates drop-down menu shows supplied templates, and any custom templates you may have added. These templates define the look and feel, arrangement, orientation, and so on, of the report output. To include the start time, end time, scan limit, device group, storage group, and devices information (used to run a report) in a report, choose the "BlankWithHeader" template.</p> <p>See "Applying Report Template Styles" on page 232 for more information on working with templates.</p>
Report Format	<p>Select the default format for the report.</p> <p>For information on available formats, see "Report File Formats" on page 179.</p>
Report Contents	Select whether report should detailed or summarized.

Select Display Fields (Base Query and Fields)

Each report is built on a base query. Available queries are provided in the drop-down menu on top of the "Select Display Fields" box.

For instructions on how to view a list of the default search fields, see ["Viewing Default Fields" on page 324](#). For information about custom schema fields added to the default schema, see ["Adding or Importing Schema Fields" on page 345](#).

When you select a query, the data fields it contains are shown in the Available Fields list. You can select which data fields you want to use in your report, or use them all. You can edit the selected query by clicking on the **Query Editor** link. (For information on building new queries, see ["Setting up Queries" on page 201](#).)









In addition to the fields in the WHERE clause of the query, the fields in the SELECT clause also need to be indexed to yield faster report generation. For more information about indexing fields, see ["Indexing" on page 121](#).

Enter a meaningful title for the report in the **Report Title** field and select whether the report contents should be Detailed or Summarized in the **Report Contents** field. The report title is the text that appears as the title on top of a report.

Select the query you want to use for the report from the drop-down list located on top of the Select Display fields section. The Available Fields list is populated with the fields defined in the selected query.

Select the fields to use in the report by moving fields from Available Fields into the Selected Fields list.

- Select a field in Available Fields and click  to move it into the Selected Fields list, or click  to add all fields.
- To “de-select” fields that you do not want in the report, select a field in the Selected Fields list and click  to move it back to the Available Fields list, or click  to “de-select” all fields.
- Use the move up  and move down  arrows to order the Selected Fields.



For information on how to create query objects for use in reports, see [“Setting up Queries” on page 201](#). All available queries, including new queries you create, show up in the drop-down menu in the Select Display Fields section of the Adhoc Report Designer.

Select Filter Criteria

Figure 5-18 Report Filter Criteria

Filter criteria defined as part of a report design is built in and saved with the report. When other users run the report, they will get the built-in filters by default. You can also set filter criteria and row limits on an ad-hoc basis when you run a report. However, values set at run time are not built in to the report like those set at design time. Run-time parameters are only applicable to a particular report run and do not persist.



If a report does include default filter criteria, users have the option to run the report with the defaults, or modify or remove the built-in filters at run time. For more information, see [“Run Report Parameters” on page 178](#).

You can set filters on the results of the base query with logical expressions to narrow the focus of the report results. For example, you could set the filter criteria on a report on Top Password Changes to report only on password changes related to specified usernames or involving specified IP addresses. You can limit the number of rows in a report by defining a Max. Rows value.

Query designers can build in “mandatory filtering” on a specified field or on “any” field, which requires filtering on one or more fields of your choice. If the query you choose for this report has mandatory filtering, the “Select Filter Criteria” panel title and one or more

fields are with a red asterisk. For more about mandatory filtering, see [“Mandatory Filtering” on page 209](#) under [“Setting up Queries” on page 201](#).

Table 5-11 Select Filter Criteria Options

Option	Description
Maximum Rows (Max. Rows)	<p>Specify the maximum number of rows in the report output. Results that push the number of rows beyond the Max. Rows limit you define will not be included in the report.</p> <ul style="list-style-type: none"> If you select set Max. Rows and also specify grouping under Set Grouping (as described in “Select Grouping” on page 191), you may get a different result than if you just specified Max. Rows without grouping. Setting this field to 0 returns an unlimited number of rows. Increasing the maximum rows for report may not always increase the number of rows returned by the report. If the query invoked by the report limits the number of rows returned, increasing the Max. Rows setting in the report has no affect. For example, if you edit the NIST IR Top 10 High Risk Events report and change the value in the Max. Rows column from 10 to 20, when the report is run report only 10 rows are returned. This is because the query invoked by the report is returning 10 rows. However, you can limit the number of rows returned by the report to a number less than the default value. For example, if the value of the Max. Rows field is changed from 10 to 5 for the NIST IR Top 10 High Risk Events report, this report returns 5 rows during run time. You can increase the number of rows returned by editing the query and changing the number of rows returned by the query and change the number specified in the Max. Rows field of the report.
Field	<p>The Fields will be populated with event data fields specified in the base query. (Fields will generally equate to columns in reports.)</p> <p>Select a field on which to filter.</p> <p>To add another filter (“Field” on which to filter), click  (Add Filter).</p> <p>To remove a filter, click  (Remove Filter).</p> <p>For instructions on how to view a list of the default search fields, see “Viewing Default Fields” on page 324. For information about custom schema fields added to the default schema, see “Adding or Importing Schema Fields” on page 345.</p> <ul style="list-style-type: none"> Multiple filters with conditions set on different fields will be AND’ed together. Multiple filters with conditions set on the same field will be OR’ed together. <p>For example, if you want to filter on events to return data based on a value/count (of rows or other) between 90 and 100, use the Between criteria to do this (for example, <Field> Between 90 and 100)</p> <p>Setting two filters on the same field with criteria “Above 90” and the other as “Below 90” would not give you the data you are looking for. Only one of these filters would be triggered.</p> <ul style="list-style-type: none"> If the query you choose for this report has mandatory filtering, the “Select Filter Criteria” panel title and one or more fields are marked with a red asterisk. For more about mandatory filtering, see “Mandatory Filtering” on page 209 under “Setting up Queries” on page 201.
Criteria	<p>Select a logical operator. (For example, Is, Is Not, Starts With, Ends With, Contains, and so forth.)</p> <p>Note: To make the query case-sensitive, select the Match Case option for your operator.</p>
Value	Select a value to complete the conditional filter expression.

Select Grouping

Figure 5-19 Grouping Items by Field in a Report

Define group requirements to arrange the report information into logical groups based on particular fields. You can create multiple groupings for report results.

For example, if the report uses a query that includes a Date field, you can group results by date. You could add additional statements to group by “User Name”, “Source Address”, “Destination Address”, and so forth, depending on what other fields are available in the report query.



- If you select set Max. Rows under **Select Filter Criteria** (as described in [“Select Filter Criteria” on page 189](#)) and also specify grouping, you may get a different result than if you just specified Max. Rows without grouping.
- A report that has a group defined can only display up to 100,000 lines.

To define a group:

- 1 Select a field by which you want to group (as described in [Table 5-20 on page 236](#)).
- 2 Select the order of arrangement of group (as described in [Table 5-20 on page 236](#)).

Table 5-12 Select Grouping Options

Option	Description
Field	<p>The Fields will be populated with event data fields specified in the base query.</p> <p>Select a field by which to create a group.</p> <p>To add another field for a grouping, click (Add Group).</p> <p>To remove a group-by field, click (Remove Group).</p>
Order	<p>Select the order of arrangement of group:</p> <ul style="list-style-type: none"> • Ascending • Descending
Ranking Field	
Ranking Function	

- 3 If you want to set sub-groups, specify details in the “Then By” fields. For example, if your report uses a query that reports on password changes and includes a “User Name” field, you might want to sub-group the results for each date by “User Name”.

Use the (Add Group) and (Remove Group) buttons to add or remove “Then By” fields for sub-groups.

The report will generate records organized and grouped in the specified order.



Alternatively, you can specify only a sort order (instead of groups). See also, [“Sort Order” on page 192](#).

Select Totals

Figure 5-20 Showing Totals on Fields in a Report

You can specify the summary (total) fields. You can apply a summary on any of the following levels:

- Report
- Page
- Group

To specify summary details:

- 1 From **Field**, select the field that will be processed to calculate summary information.
- 2 On the same row, from **Function**, select the summary function.
- 3 On the same row, from **Level**, select the level at which you want the summary.



If a Total is applied to a field that is not already in the Selected Fields list, that field is automatically added to the Selected Fields list.

Sort Order

In case you do not want “grouped” report results (as described in a [“Select Grouping” on page 191](#)), but you do expect “sorted” results, then specify a Sort Order (instead of grouping).



A report that has a sort order defined can only display up to 100,000 lines.

Figure 5-21 Sort Order for Items in a Report

You can have up to three levels of sorting.

To specify a sort order:

- 1 In **Field** (on the right of Sort By), select the field on which you want to sort the report.
- 2 In **Criteria** (in the same row), select the sort criteria.
- 3 Repeat [Step 1](#) and [Step 2](#) by providing values in the “Then By” rows to specify more sorting criteria.

Highlighting

A report can include multiple levels of “highlighting” for specified fields. Highlighted items can serve as visual alerts on generated reports when specified set conditions are satisfied.

Figure 5-22 Highlighting Items in a Report

To set up a highlight:

- 1 In **Highlight**, select the field that should be highlighted. Select Entire Row to highlight entire record.
- 2 In **Using Style**, select the style to be applied to highlight it.
- 3 Select **Alert** checkbox to receive a visual alert on report viewer.
- 4 In **Field**, select the fields which will be evaluated for highlight (alert).
- 5 In **Level**, select the level at which the selected field should be evaluated:
 - ◆ DETAIL evaluates each row (record)
 - ◆ REPORT evaluates at the end of report
 - ◆ Respective groups evaluate at the end of each group
 - ◆ PAGE evaluates at the end of the page
- 6 When REPORT or PAGE is selected in Level, select a Function to be applied.
- 7 Select **Criteria** and specify its **Value**.

Click (Remove Condition) on the left of the criteria entry to delete an entry. Click (Add Condition) to add another entry.

Create Matrix

You might choose to include a matrix in your report, since it presents a summary of data. Make sure that the appropriate query object is selected (under “Select Display Fields”).

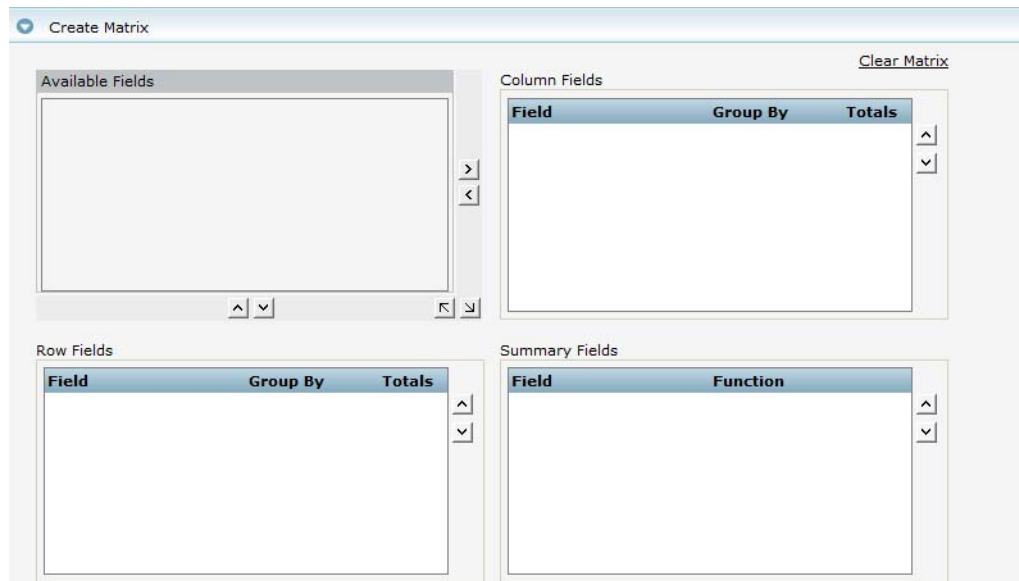


Figure 5-23 Adding a Matrix to a Report

To create a matrix:

- 1** To place a field in Row or Column, click the field and drag it to the **Row Fields** or **Column Fields** boxes.
- 2** To place a field as a cell (summary), click the field and drag into the **Summary Fields** box.
- 3** Select a **Function** from the drop-down menu provided for a field placed in **Summary Fields**.
- 4** Optionally, for numeric or date fields in columns or rows, specify a **Group By** function in the drop-down.
- 5** Optionally, for fields in columns or rows, check **Totals** checkbox to get total row / column.

Select a field and click to add that field to the matrix as one of the **Column Fields**.

Select a field in Column Fields and click to remove it from the matrix.

Select a field and click to add that field to the matrix as one of the **Row Fields**. Select a field in Row Fields and click to remove it from the matrix.

Select a field and click to add that field to the matrix as one of the **Summary Fields**. Select a field in Summary Fields and click to remove it from the matrix.

To move a field up or down, select the field and click (Move up) or (Move down), to move the field in the respective direction.

To remove all settings and contents of the current matrix, click **Clear Matrix**.

Create Chart

For pictorial representation of summary data, you can add a chart on your report. Make sure that the appropriate query object is selected (under “Select Display Fields”).

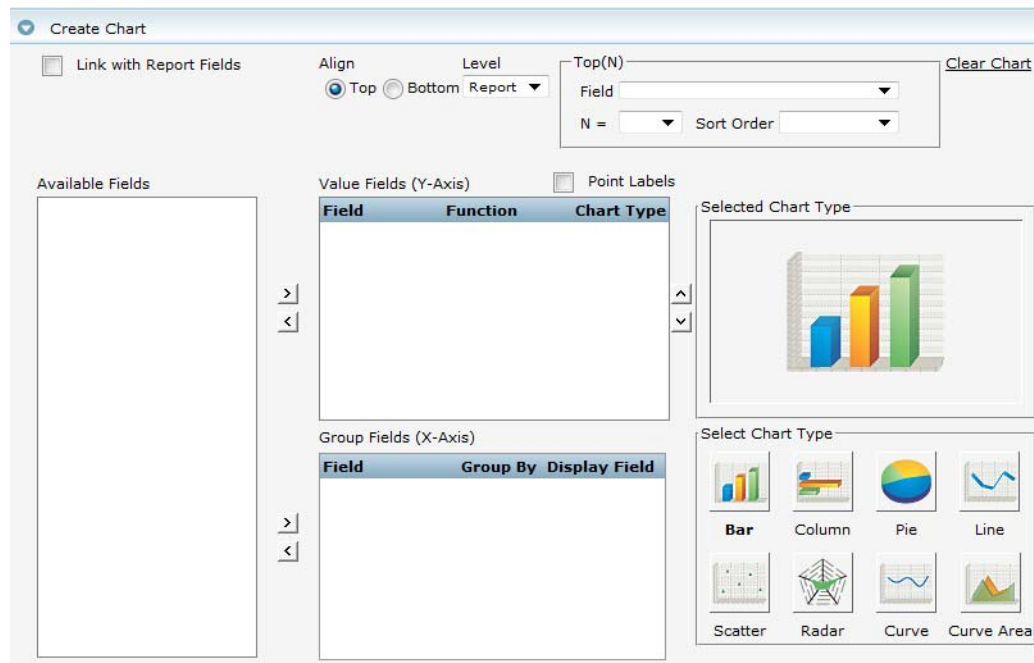


Figure 5-24 Adding a Chart to a Report

For pictorial representation of summary data, you may choose to have a chart on your report. Make sure that the right query object is selected (under Select Display Fields).

Chart Placement

Chart Placement is important when the chart is placed on the report along with other component. Specify chart placement preference using the Align option:

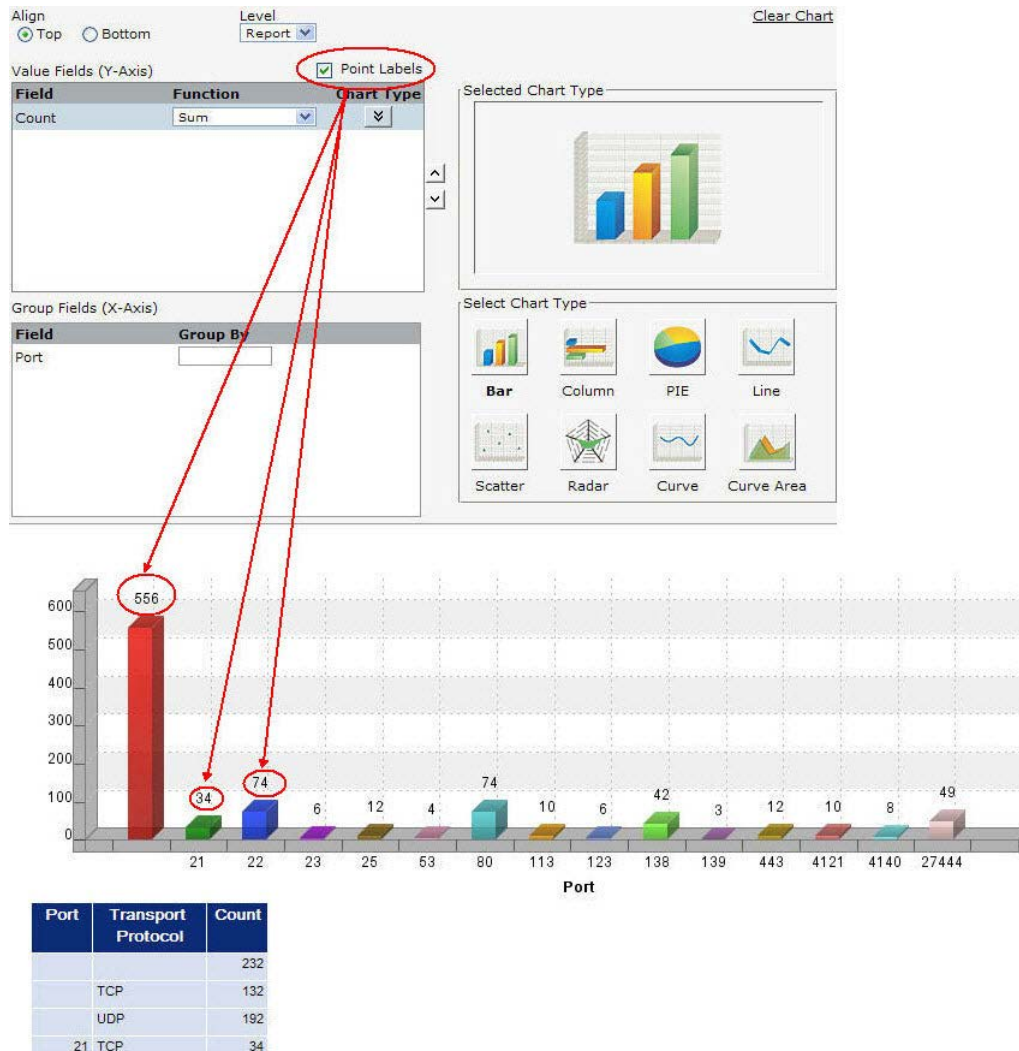
- Select **Top** to place the chart above other components.
- Select **Bottom** to place the chart below other components.
- In **Level**, select PAGE to plot chart having page level data. Select REPORT to plot chart from data that has come from entire report.
- Top(N)
 - ◆ Field
 - ◆ N=
 - ◆ Sort Order

Chart Type

Select the chart type by clicking button (image) from **Select Chart Type** area. The image corresponding to the chart you select is displayed in the **Selected Chart Type** box at the top.

Select Point Labels

Select this setting to show the number of matches for a value of a field in a chart, as shown in the following figure.



Set Value Fields (Y-Axis)

- 1 Click and drag the Field in **Value Fields (Y-Axis)** box, or use the **>** button (Add field) to add the selected field.
- 2 Select summary function for the field.
- 3 To select a different chart type, click the button on the right to open a box with chart types. Select the type you need. Follow steps 1 through 3 above for each attribute to be placed as series. To re-position fields, select a field and click **^** (Move up) or **v** (Move down) as needed.

Set Group Fields (X-Axis)


- 1 Click and drag the field in **Group Fields (Y-Axis)** box, or use the  button (Add field) to add the selected field.

- 2 Select the method to group (for Numeric or date type).

You can specify groups in numeric fields. For example, to have groups of 10, specify 10 in Groups box.

You can specify groups in date fields. From the drop-down box select from Day, Week (Sunday to Saturday), Month, Quarter (Jan-Mar, Apr - Jun, Jul - Sep, Oct - Dec), Year.



To remove fields from Value fields (Y-Axis) or Group Fields (X-Axis), drag them out of the respective box or use the  button (Remove field) on selected fields.

To remove all settings and contents of the current chart, click **Clear Chart**.

Editing a Report

You can use the Report Designer to edit existing user-designed reports. (The supplied reports are not editable.)

To edit an existing report:

- 1 From any Report list in the Report Explorer, select the report and click the **Customize Report** button for the report you want to edit.

This brings up the Report Designer for the selected report.

- 2 Modify the report as needed (via the settings described in [“Creating New Reports” on page 186](#)).
- 3 (Optional) Before saving the report, you can run it to ensure that the changes you expected in the report output suit your needs. To do so, click **Run**. (For more information see, [“Adhoc Report Designer” on page 197](#)).
- 4 Click **Save**.

See also [“Quick Start: Base a New Report on an Existing One” on page 186](#).

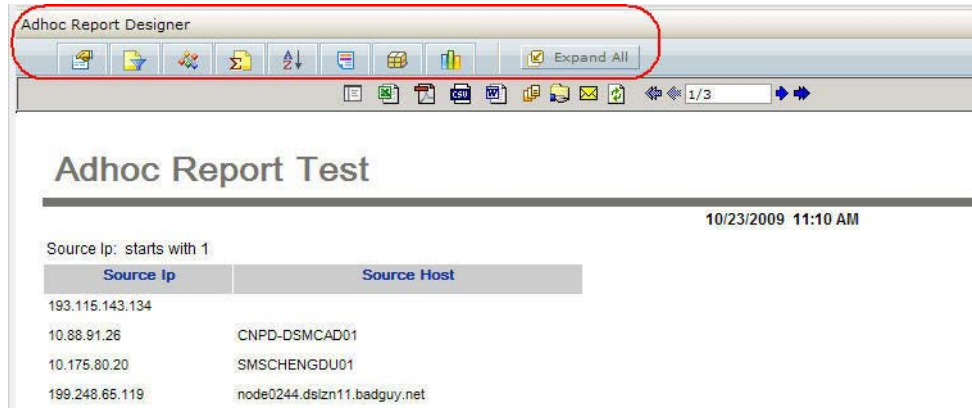
Private Reports

If you have access rights to “view, run, and schedule all reports”, you can create *private* reports. If you do not have permissions to edit a *public* report that you want to modify but you do have permissions to create private reports, then you can save the public report as a private one and edit the private report.

For more about publishing a report as “public” or “private”, see [Table 5-8 on page 181](#). For more information on setting permissions, see [“Setting Access Rights on Reports” on page 198](#)









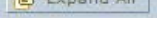
Adhoc Report Designer

Once you edit a report, you can run it before saving it to ensure that the report output is as you expected. When you run a report in this fashion, an Adhoc Report Designer menu bar is displayed at the top of the newly run (unsaved) report, as shown in the following figure.



The Adhoc Report Designer is useful in adding formatting and display elements to a report definition and viewing the output with those elements before saving the report definition. For example, you can specify a sort pattern or add a chart to a report.

The following table lists the various options available in the Adhoc Report Designer menu bar.

Menu Option	Description
	Select display fields See "Select Display Fields (Base Query and Fields)" on page 188 for more information.
	Specify filter criteria See "Select Filter Criteria" on page 189 for more information.
	Specify grouping See "Select Grouping" on page 191 for more information.
	Specify the summary (total) fields See "Select Totals" on page 192 for more information.
	Specify sort order See "Sort Order" on page 192 for more information.
	Set up highlighting See "Highlighting" on page 193 for more information.
	Include a matrix See "Create Matrix" on page 194 for more information.
	Create a chart See "Create Chart" on page 195 for more information.
	Expand all of the above listed menu options.

Setting Access Rights on Reports

Administrators can set access rights on various report categories, reports, and report options (view, publish, edit, and so on) based on user roles and Logger Report Group affiliation. For example, you can grant users privileges to view some reports but not others, to view but not schedule or publish a report, or to view and schedule but not edit a report.

(This is also noted with regard to user perspective at [“Task Options on Available Reports” on page 173.](#))

Access rights are given at the folder level. If you want to give access only to specific reports, you can put them in their own folder and give access to it. Access rights on report options are configured and managed with the User/Groups option on the Logger System Admin page. For more information on System Admin User/Group management, see [“User Management” on page 402.](#)

Determining What Access Rights to Give a Group or User

When setting access rights for a user or group, be sure to give the user all the necessary permissions. In order to access a particular child node, users need access rights to all higher nodes in that branch of the tree.

To determine the necessary rights for a report, open the report tree to that report.



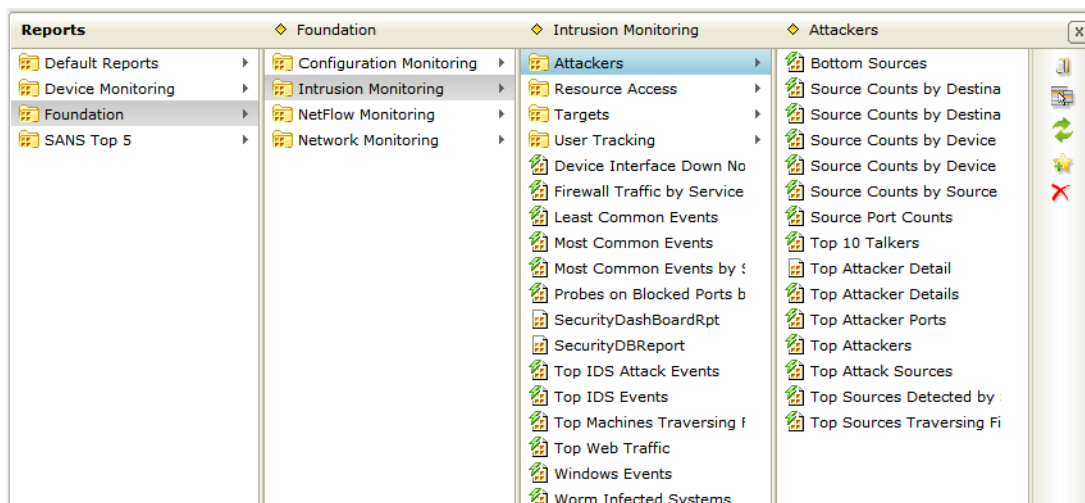
The branches of the tree and give the user or user group access to all.

Example: Giving a User Group Access Rights for a Report

Suppose you want to give a Group the rights to view, run, and schedule, but not to change the Attackers reports. To determine the necessary rights, scan the report tree and note the nodes.

To view Attackers Report tree and determine the necessary rights to access it:

- 1 Click **Reports** from the top-level menu bar.
- 2 Click **Report Explorer** in the **Navigation** section on the left panel.
- 3 In the Reports screen, navigate to the group of reports you want to give access to. For the example, click **Foundation > Intrusion Monitoring > Attackers**.



- 4 Make a note of each node you open.

Now that you know the nodes you need to give access rights to, you can set them from the System Admin menu.

To create a new User Group and give it Logger Reports Rights:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section on the left panel.
- 3 Open the **Groups** tab, and click **Add**.
- 4 Type in a Name for the group and add a description.
- 5 Select **Logger Reports** from the Group Type drop down menu.
- 6 Click the arrow to display the list of Logger Reports Rights.
- 7 Click **Clear All** to remove all permissions.
- 8 Click the box next to each permission you want to give the user group.


For the example, you noted Foundation, Intrusion Monitoring, and Attackers, and you wanted to give the rights to view, run, and schedule these reports. Therefore, put a mark in the box next to each of the following access rights:

Report folder [Attackers]: view, run, and schedule reports

Report folder [Foundation]: view, run, and schedule reports

Report folder [Intrusion Monitoring]: view, run, and schedule reports

- 9 Click **Save and Edit Membership**.
- 10 Click **Add** in the Edit Group Membership screen.
- 11 Put a mark in the box for the user you want to add to the group, and click **OK**.
- 12 Log in as a member of the group you created and test whether you can perform the desired functions. For the example, the user should be able to view, run, and schedule the Attackers reports only.



Setting up Queries

Query objects are queries (along with additional metadata) designed and stored as a part of the Logger Reporting suite on the Report. Query objects are used as the basis for designing reports.



Note

Some queries may require parameters. We recommend first designing all needed parameter objects before creating the query object that will use those parameter objects.


For information on developing parameter objects, see [“Working with Parameters” on page 223](#).

To view and work with Logger Report queries, click Design | **Queries** on the Reports left pane. The contents for the selected query is displayed. To view the contents of a different query, click the **Open** button and expand the category under which the query is stored then select a query from the list and click the **Select Entity** button.

Logger Reporting provides a set of pre-built queries, which are used as the basis for the system defined Reports and Solutions Reports to address common security use cases (as described in [“Explorers” on page 139](#)).

For instructions on how to view a list of the default search fields, see [“Viewing Default Fields” on page 324](#). For information about custom schema fields added to the default schema, see [“Adding or Importing Schema Fields” on page 345](#).

You can use a provided query object “as-is” as the basis for your own reports, or design new query objects on the Query Object List page. You can use existing query objects as a starting point for new ones. You can search for an existing query, as shown in the following figure.



To do so, either

- Enter the first few letters with which the query name begins (if the “Starts With” search criteria is selected) in the text box above the list of existing queries, OR

- Enter a word or part of a word that the query name contains (if the “Contains” search criteria is selected) in the text box above the list of existing queries.



Modifications to reports and other ArcSight-defined content may be overwritten without warning when the content is upgraded. It is not good practice to modify ArcSight-defined content directly.

Make modifications to a copy of any ArcSight-defined content as a general practice, and subsequent upgrades will not affect the modifications.

This topic explains how to design new query objects (either from scratch or based on existing ones).

How Search and Report Queries Differ

Even though a search and a report query perform the same function—finding events that match specific conditions—the two queries are distinct in these ways:

- You use Logger’s in-built SQL Editor to create a report query in SQL. (The SQL Editor automatically checks the syntax of the query before running it.)
- You use the Logger’s Search UI to create a search query. The query can be specified using plain English keywords, field names, or regular expressions. See [“Searching for Events on Logger” on page 106](#) for more information.

However, report queries and field name queries can utilize indexed fields to expedite the underlying search.

Overview of Query Design Elements

To create a new query object, you need to specify a query name, define the SQL logic, and save it. The data source for Logger Report queries is always the Logger database(s), so there is no need to specify this as part of the query object.

Optionally, you can specify formulas, set field properties, define formatting (look-and-feel), define field groups, provide hyperlinking, define lookup values, and build mandatory filtering into the query.

Creating a Copy of an Existing Query

To use an existing query object as the basis for a new one, copy the query object you want to start with as follows:

- 1 In the **Query Explorer**, click on a category and select the name of the query that you want to copy from the query list.

To search for an existing query, do one of the following:

- ◆ Enter the first few letters with which the query name begins (if the “Starts With” search criteria is selected) in the text box above the list of existing queries.
- ◆ Enter a word or part of a word that the query name contains (if the “Contains” search criteria is selected) in the text box above the list of existing queries.

- 2 Click the **Copy Query Object** button, then click the category name under which you want to place the copied query, and click the **Paste** button.

A temporary version of the new query object is created with the same contents as the original and the same name pre-fixed with “Copy of”.

- 3 Modify the query name by editing it in the **Name** field (unless you want to keep the default name of “Copy of <OriginalQueryName>” for now).

Designing a New SQL Query

When using the Query Editor, be sure to use the appropriate SQL syntax for your data type. For example, to call a string data type, you must enclose the string with single quotes, as in the query below.

```
select arc_deviceVendor from events where lower(arc_deviceVendor) = 'arcsight'
```

For instructions on how to view a list of the default search fields, see [“Viewing Default Fields” on page 324](#). For information about custom schema fields added to the default schema, see [“Adding or Importing Schema Fields” on page 345](#).

To design a query using the Query Explorer:

- 1 In the Query Explorer, click the **Create Query Object** button. The Query Object Editor opens.
- 2 In **Name** field, specify a unique name for this query object.
- 3 Under **SQL**, click **Edit** to design the SQL statement.

The SQL Editor loads in a new window by default, which is generally preferable because it allows you to view both the main Query Object List page (query editor) and the SQL Editor at the same time. (If you want the SQL to load in the same window, click to uncheck this option before clicking the Edit button.)

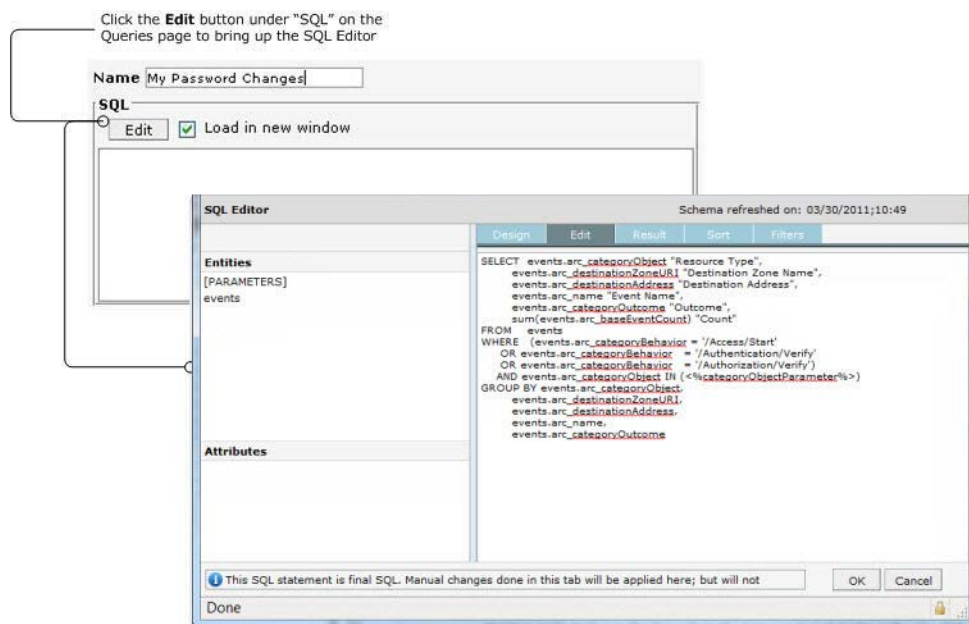


Figure 5-25 Query SQL Editor

- 4 Use the SQL Editor to define the query statement. (See [“Defining SQL in the Editor” on page 215](#).) Report queries are case insensitive.
- 5 Click **OK** to temporarily save the SQL statement for the query.

The SQL you defined is displayed in the SQL box on the main Query Object List page.

Similarly, any fields you defined in the SQL Editor are displayed in the Fields list on the Query Object List page.

- 6 Click **Save** to save your work as part of the query object.



You must click **Save** on the main Query Object List page to save updates made in the SQL Editor as part of the query. If you navigate away from this page without clicking Save, edits you made in the SQL Editor since the previous Save will be lost.

Field Attributes and Properties

To set Field attributes, select a field under **Fields** and edit the properties associated with that field.

Figure 5-26 Query Field Attributes


You can set the following properties on fields in a query.

Table 5-13 Query Field Attributes

Option	Description
Field	Name of field (as received from data source).
Caption	The text that will appear as a caption when this field is selected for placement on the report.
Width	Number of characters for the selected field.
Align	Sets alignment for the selected field.
Hidden	Hides the associated field so that it is not available to be placed on report. This field will also not be available for sorting as well as filtering.
Data Type	Sets the data type for field from Date, Character, or Number. This is especially useful when field selected is XML type data source and you need to set it as number or date. Similarly, when a field that is character (having numeric value) is supposed to be used in calculation.

Specifying Output Format for a Field

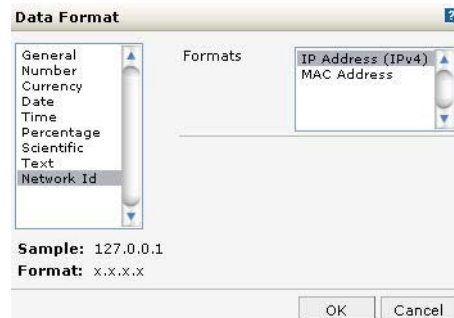
If you specify the output format for a query field here, at run-time, the report output will adhere to the specified formatting.

- 1 From Fields list, click (select) the field for which you want to define an output format. (The selected field is bold.)
- 2 Click  button next to the Output Format field to launch the Data Format dialog.

- 3 Select the appropriate format and provide necessary values for that format.



The default date/time in reports does not include the time of day. You must choose a date format that includes HH:MM:SS to include the time.




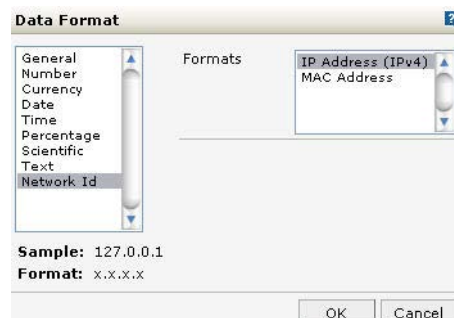
- 4 Click **OK** to accept the changes and close the dialog.

Characters for the selected format are displayed in the Output Format entry field.

Specifying Input Format for a Field




If you specify the input format for a query field here, at run-time the report containing this query will accept data only in the format specified.

- 1 From Fields list, click (select) the field for which you want to define an input format. (The selected field is bold.)
- 2 Click  button next to the Input Format field to launch the Data Format dialog.



- 3 Select the appropriate format and provide necessary values for that format.
- 4 Click **OK** to accept the changes and close the dialog.

Characters for the selected format are displayed in the Input Format entry field.

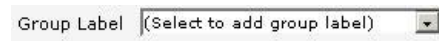
Format			
Width	<input type="text" value="15"/>	Output Format	<input type="text" value="x.x.x.x"/> 
Align	<input type="text" value="Right"/> 	Input Format	<input type="text" value="x.x.x.x"/> 

Grouping Fields

When fields in a query object are grouped, they are displayed within a group header in the Report Designer. All fields in the group can be selected or removed from the report with a single click. Once groups are created, fields can be assigned to groups.

To create groups:

- 1 In Group Label drop-down box, double-click the **Select to add group label** option.



Group Label (Select to add group label)


- 2 Enter a group name.
- 3 To create more groups, repeat [Step 1](#) and [Step 2](#).

To assign fields to a group:

- 1 From the Fields list, select the field.
- 2 From the Group Label drop-down box, select a group.

The selected field will be part of that group.

To remove a group:


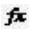
- 1 Select the group name in the Fields list.
This automatically populates the Group Label field with the selected group name.
- 2 Click  (remove button) next to the Group Label field to remove the selected group.

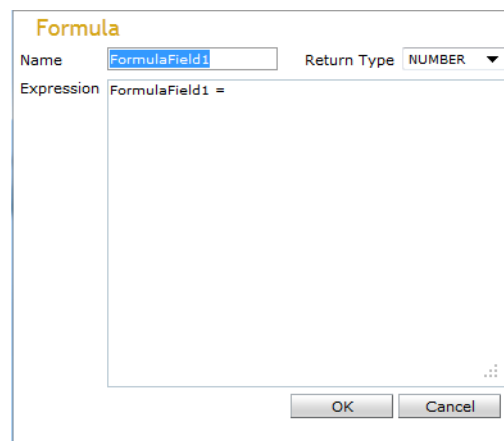
Formula Fields

Formula fields are custom fields you create to address particular scenarios during report processing. In a business finance application, a formula field might be used to determine “gross salary” or “grand total”. Formula fields and their values are not stored in the Logger Report server database; but rather are used during report processing and discarded once the report is generated.

You can embed formula fields query objects. A formula field can include any field or formula available to a query object.

To create a formula:

- 1 Click  (Add New) under the  tool palette (next to the Fields list) to get the Formula dialog.
- 2 Specify the formula and click **OK** to save the formula and close the dialog. (See [“Syntax for Formulas” on page 207.](#))



Formula

Name: FormulaField1 Return Type: NUMBER

Expression: FormulaField1 =

OK Cancel

- ◆ In the **Name** field, specify a unique name to identify the formula.

- ◆ From the **Return Type** drop-down menu, select the type of the value the expression returns. (NUMBER, CHAR, DATE, or BOOLEAN)
 - ◆ In Details area, specify the formula. (See [“Syntax for Formulas” on page 207.](#))
- 3 Click **OK** to save the work and close the dialog box.

The new formula is listed in the Fields list. Formula names shown in the list are pre-fixed by **fx** to indicate they are formulas.

Positioning Formulas in Fields List

Select a formula and click (Move up) or (Move down) as needed to shift the position of the selected formula in the list.



Formulas further down in the list can use the formulas above them.

Avoid the opposite; formulas higher in the list should not use the formulas below them.

Syntax for Formulas

The general syntax for formula is:

```
FormulaName = formula
```

where, FormulaName is the same as specified in the **Name** field on the Formula dialog.

In general, use JavaScript syntax to create formulas.

A formula can include:

- Field names
- Variables (custom or supplied)
- “if” and “nested if” constructs
- logical operators

For formulas that contain multiple statements, use a semicolon “;” as a separator between two statements.

Examples

```
NewForm1 = var a = 5 ; b = 3 ; if (a!=b) { f = a } {NewForm1=f}
```

```
TotalAmount = var total ; if (unitprice < 10 ) {total = unitprice*quantity} else {total = unitprice} {TotalAmount = total}
```

Importing Field Attributes

You can import field attributes from other Logger Report queries and apply the imported attributes to the currently selected field in your query. Leveraging attributes from existing queries can save time and re-work, and also serves as a learning tool.

You can import the following field attributes from one query into another:

- Captions
- Format (including Width, Alignment, Input, and Output formats)
- Data Types
- Hidden properties
- Group Labels
- Hyperlinks
- Lookup Values

You can select a field from which to import attributes from any of the saved query objects on the Logger Reporting server. Imported attributes can come from one field in another query, or from multiple fields.

To import field attributes:

- 1 Open the query in the Query Object Editor, select the query object into which you want to import field attributes (the “local” query you are editing), and click **Import** to bring up the Import Attributes dialog.

Figure 5-27 Importing Attributes from One Query into Another

- 2 From the **Source Query** drop-down menu on the Import dialog, select the query object from which you want to import field attributes (the “remote” query with the attributes you want to copy).

The Field drop-down menu is populated with the fields in the selected query.

- 3 In the **Source Field** drop-down menu, select the field in the remote query whose attributes you want to import (copy).
- 4 Select field attributes to import by clicking (checking) the checkboxes for attributes you want.
- 5 From **Target Fields** drop-down menu, select the target field in your local query to which you want to copy the attributes.



Note

For successful field attribute import, consider the data types of source and target fields. For a valid import, data types for source and target fields must generally match.

Lookup values will not import if the data type of the target field is NUMBER.

- 6 Click **Apply** to save current selections and keep the dialog open.



Caution

A field attribute import cannot be revoked. Please make sure you are importing the right attributes before you click **Apply** or **OK**.

- ◆ Click **Cancel** to abandon selections made after last Apply button and close the dialog. (Clicking Cancel will not revoke changes already applied.)
 - ◆ Click **OK** to save (apply) current selections and close the dialog.
- 7 To import selected field attributes to another target field, repeat these steps with a different target field selected.

To select from different fields in the same query, or different queries, choose different options for **Query** and **Field** at the top of the dialog.

Mandatory Filtering

You can provide built-in filters for a query when you want users to apply one or more filters when designing and running reports that use that query. Building in mandatory filtering at the query level can save unnecessary data transfer from the server database during report run time.


You can configure mandatory filtering in either of these ways:

- Mandating filtering on *any field*. Report designers can decide which field to filter on at report design time.
- Mandating filtering on a *specific field*. Report designers are required to filter on the specified fields at report design time.

To configure a query for mandatory filtering:

- 1 Select (check) the **Mandatory Filtering** checkbox to enable mandatory filtering.

- 2 To specify a field for mandatory filtering, choose the field you want from the **On Field** drop-down menu. If you do not want to specify a field for mandatory filtering now, leave it as **Any**.

- 3 Click  (Add Filter) to get another row for mandatory filtering, and repeat [Step 2](#) above.

To remove a field filter:

To remove a mandatory filter field, click  (remove button) next to the row you want to remove.

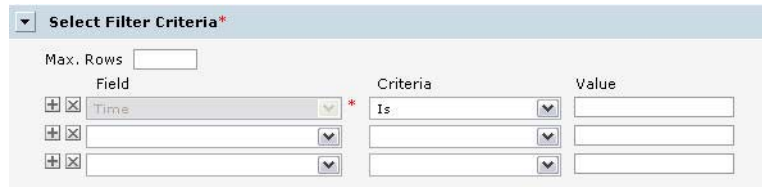
To disable mandatory filters:

To disable mandatory filters (but not remove the specified fields), uncheck the **Mandatory Filtering** checkbox. (Click on it if it is enabled to toggle it off.)

Effect of Mandatory Filtering on Report Design

Mandatory filtering comes into play during report design time with regard to selecting filter criteria. (See [“Select Filter Criteria” on page 189](#) under [“Designing Reports” on page 185](#).)

When a user working with the Report Designer to create/edit a report selects a query object (data source) that has a mandatory filter, both the “Select Filter Criteria” panel title and the relevant fields are marked with a red asterisk.



Field	Criteria	Value
Time*	Is	

Figure 5-28 Mandatory Filtering on a Field Shown in the Report Designer

The Report Designer “Select Filter Criteria” panel includes one row for each field configured for mandatory filtering in the base query (all marked with red asterisks).

For each mandatory field configured with a *specified field* in the base query, a named field is provided in the Report Designer “Select Filter Criteria” with its drop-down menu grayed out (disabled). This requires the report designer to build the report so that it filters on the specified field.


For each mandatory field with “Any” (*any field*) as the selected value in the base query, a “blank” field is provided in the Report Designer “Select Filter Criteria” with its drop-down menu enabled. In this scenario, the report designer is required to build the report to filter on a field, but it can be any field provided by the query to the report via the filter criteria drop-down list.

So during report design, filters must be provided for all the rows marked with red asterisks, but mandatory filtering on “any” field gives the report designer a little more leeway than mandatory filtering on a specified field.

Specifying a Hyperlink on a Field

You can make a field a clickable hyperlink that links to a specified URL or report. A report based on a query with hyperlinked field(s) will provide links to intranet or external Web pages and/or “drill-down” reports.

To make a field a hyperlink:

- 1 From Fields list in the query, click (select) the field you want to be the hyperlink. (The selected field is bold.)
- 2 Click the ellipsis  button next to the Hyperlink option to launch the Hyperlink Options dialog.






Figure 5-29 Making a Field in a Report a Hyperlink

- 3 Depending on the type of hyperlink needed, select either **URL** or **Drill Down...**, and specify values appropriately.

Link Type	Settings
URL	<ul style="list-style-type: none"> • Select URL • Provide the link address in the box URL box (HTTP or HTTPS address, file path, etc.) • Choose a Target window or frame from the drop-down menu, depending on how you want the URL target to be displayed (same window, new window, and so on).

Link Type	Settings
Drill Down to Another Report	<ul style="list-style-type: none"> Select Drill Down to another Report Choose a Target window or frame, depending on how you want the new report to display <p>Note: A report may have mandatory parameters. If the value of a mandatory parameter is not specified, the report run may fail, generate errors, or provide invalid results.</p> <ul style="list-style-type: none"> If the target report needs system parameters to run, specify these along with associated values. Add and remove rows in the same way as for report parameters. For details, see “System Parameters and Associated Values” on page 212. <p>Even if the target report (the report you are linking to) does not need any report parameters to run, specify the following parameter in the Report Parameter section. This parameter is required for the drill down functionality in a report to work:</p> <p>Report Parameter: REQ_SD Value Field: <%REQ_SD%></p> <p>Click  to add a row or  to delete a row in the Report Parameter section.</p>

- 4 Click **OK** to accept the changes and close the dialog.

The Hyperlink option for the selected field is now blue to indicate that the field is a link. (Query “Fields” list that are hyperlinks always show a blue Hyperlink option when they are selected in the list.)




Figure 5-30 Hyperlink Options

System Parameters and Associated Values

You can set the following system parameters to further specify how a target (hyperlinked) report is run and published.

Parameter	Description and Values
Priority	<ul style="list-style-type: none"> Low Medium High
Report Format	<ul style="list-style-type: none"> Choose SYS_REPORT_FORMAT to use the format of the report specified where the target report is run Or choose one of the other formats on the drop-down (described in “Report File Formats” on page 179)
Report Connection Name	<p>Report type and database.</p> <p>We recommend leaving this set to Default.</p>

Parameter	Description and Values
Save File Name	Provide a file name to be used for the target report if the report is published as an implicit operation.
Implicit Operation	Publish is the recommended default option.
Refresh Data	<ul style="list-style-type: none"> Select True to run report with latest data. Select False to run report with cached data
Prefetch Drilldown	<ul style="list-style-type: none"> Select True to enable “prefetch drill-down” and generate hyperlinked report at run time, even if user has not clicked the hyperlink in the source report. Select False to disable prefetch drill-down
Pagination	Select a pagination option for the target report: <ul style="list-style-type: none"> Single Page increases page width and length to any size Multiple Page divides in width, divide in length as per need Horizontal Breaks divides in length only, increase width to any size Vertical Breaks divides in width, increase length to any size
Show HTML Toolbar	If the target report is published or viewed in HTML: <ul style="list-style-type: none"> Set Yes to have HTML Toolbar Set No to forego the toolbar Set Multipage to provide toolbar only if report extends to more than one page.

Lookup Values for Text Fields

Lookup values are used to set a filter at report design time as well as run time.

Query objects are generally used by report designers. Query designers can configure lookup values for fields on which report designers may decide to set filters at report design time or users may want to filter at report run time.

When a report designer sets up a filter on a field, lookup values for the field are listed in a drop-down menu. The report designer can select a value and proceed with building the report. Similarly, at run time, a dialog is displayed with the field name and lookup values listed in a drop-down menu. The query will run with the filter and specified values.

Lookup values can be defined in any of the following ways:

- Predefined, to specify static values.
- SQL, to get values from the database using SQL (used in the main query or from a query setup exclusively). This way you make sure that the user selects valid options.
- Key Field, from the table used in the main query. Specifying a key field can improve performance.

To specify predefined lookup values:

- 1 From Fields list in the query, click (select) the field to which you want to assign lookup values. (The selected field is bold.)
- 2 If not checked, select (check) the **Lookup Values** checkbox.
- 3 Click (check) **Predefined** link.

Figure 5-31 Setting Predefined Lookup Values in a Query

- 4 In **Display** field, specify the value to present to the user or report designer.
- 5 In **Value** field, specify the value to be provided when the user selects the value specified in "Display".
- 6 Click to add the value set in the list of the lookup values.
- 7 Repeat the [Step 4](#) through [Step 6](#) to add all the pre-defined lookup values.
- 8 Click **Save** to save your work.

To Specify SQL or XML Lookup Values:

- 1 From **Fields** list in the query, click (select) the field to which you want to assign lookup values. (The selected field is bold.)
- 2 If not checked, select (check) the **Lookup Values** checkbox.
- 3 Click (check) **SQL** or **XML** link.
- 4 Optionally, Check (select) the **User Defined SQL** or **User Defined XML** checkbox to specify separate SQL for getting lookup values from database.

Alternatively, keep this checkbox unchecked (clear) to get distinct values using the SQL defined for the main query object.

- 5 Optionally, check **Fetch on every use** check-box to refresh the list of values at query design time, report design time, and report run time.

Alternatively, keep this checkbox unchecked (clear) to fetch values at query design time only. Values will be placed in the query object used at report design time and report run time.

- 6 From the Display Column drop-down menu, select the column to be used to display value to the user (only when SQL/XML is user defined).
- 7 From the Value Column drop-down box, select the column to be used in the filter (only when SQL/XML is user defined).
- 8 Click **Save** to save your work.

Modifying a Query Object

Use the Query Object editor to modify existing queries.



We recommend that you not modify queries provided with Logger or add-on Solution packs. If you want to use a supplied query as a starting point for your own queries, copy them and edit the copies, as described in [“Creating a Copy of an Existing Query” on page 202](#).

To modify an existing query:

- 1 In the **Query Explorer**, click the category in the Query Objects column where you have stored the query and click the **Edit Query Details** button.
- 2 Edit the query as needed (via the settings described in [“Setting up Queries” on page 201](#)) and click **Save**.

Deleting a Query Object

You can remove custom queries, but not supplied queries provided with Logger or add-on Solution packs.

To remove a query:

In the **Query Explorer**, click the category in the Query Objects column where you have stored the query and click the **Delete** button.

Defining SQL in the Editor

Each report is built on an SQL query of the Logger databases. SQL (Structured Query Language) is an ISO based standard programming language for retrieving and updating information in a database. Logger supports SQL queries, and provides an interactive, SQL Editor in which to define SQL statements.

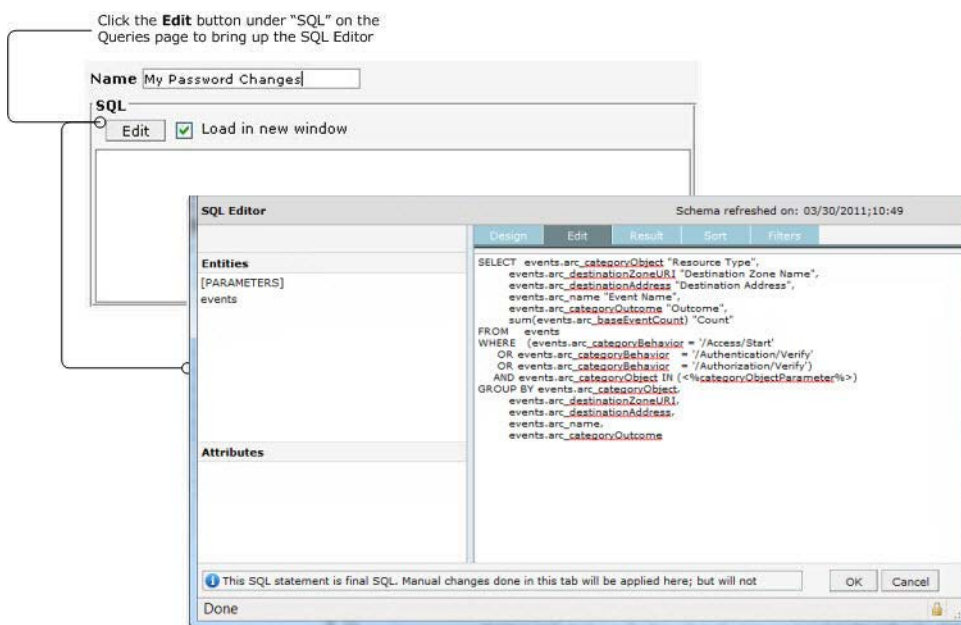


Figure 5-32 Accessing SQL Editor on the Reports | Queries page

Entities and attributes for the selected entity are listed on the left side of the SQL Editor. The right side of the SQL Editor provides tabs showing information related to the selected statement.



The Attributes list shows a few attributes that are internal to Logger. They should not be used in queries because the resulting report will not contain expected results. All attributes listed after `arc_sourceZoneResource` are internal, including `arc_eventTime`, `arc_deviceName`, `arc_rowId`, and `arc_others`.

Table 5-14 SQL Editor Tabs

Option	Description
Design	Graphical SQL query designer. Use options on this tab to design relatively simpler queries using drag and drop method.
Edit	Shows the SQL statements. A query created on the Design Tab is represented as an SQL statement on this tab. You can also write or paste and SQL directly here.
Results	Displays rows received as a result of SQL execution.
Sort	Specify sorting preferences.
Filters	Add filters to set run-time filter criteria to be included in the query.

List of Database Objects

The SQL Editor shows the **Default Connection** to the database that provides the database objects list. Logger Reporting provides a single type of object or *entity*, which is an *events* table. When you click on **events** (under Entities), event fields (attributes) are shown under **Attributes**.

Design Tab

You can design simple SQL queries on the **Design** tab using “drag-and-drop”.

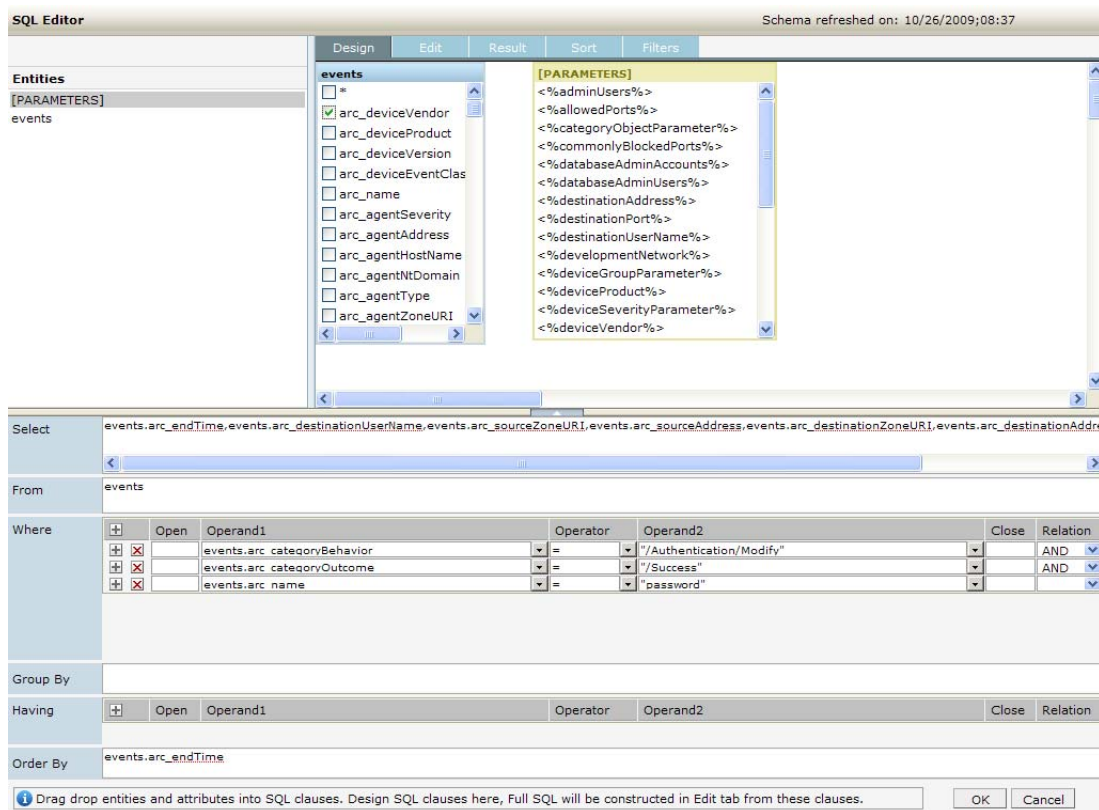


Figure 5-33 SQL Editor: Design Tab

To create an SQL query statement using the Design tab:

- 1 Under **Entities** on the left side of the editor, click **events** to select the “events” entity.
The list of event attributes is shown under **Attributes**.
- 2 Click and drag event attributes from the **Attributes** list on left side of the editor to the **Select** box on the right. The associated values are automatically displayed in the **From** clause.



The Attributes list shows a few attributes that are internal to Logger. They should not be used in queries because the resulting report will not contain expected results. All attributes listed after `arc_sourceZoneResource` are internal, including `arc_eventTime`, `arc_deviceName`, `arc_rowId`, and `arc_others`.

- 3 Repeat these steps to select other attributes from different entities.






The **events** entity must be selected (under **Entities** on the top left) in order for the event attributes to show up under **Attributes**. If no attributes are displayed, make sure you have “events” selected in the **Entities** list on the left side of the SQL Editor.

Select

The Select box shows the attributes selected for a given entity.

Where

The Where area shows the “where” clause for the query.

- To get a row at the top, click  (Insert first condition) in the left-most cell of column header.
- To get a row below current row, click  (Add a condition) in the row below which you want to add a row for condition. A row is inserted in the row below the respective row.
- To remove a condition, click  (Remove this condition) in the row for the condition you want to remove.
- To specify a where clause, form a condition by selecting Operand1, Operand2 and Operator.
- To join conditions, create two conditions, and select a relation in the right-most column of the first condition (of the two being joined).
- To group conditions, specify opening brace and closing brace in the right row.

Group By

In the Group By clause you can provide grouping criteria for the SQL statement. To place an entity in Group By, click the entity in the Entity List and drag it in the box below Group By.

Having

To build a “Having” clause, use the same settings as described in the “Where” clause. See [“Where” on page 218](#).

**Note**

Be sure to include appropriate summary function in “Select” clause so that it can be used in the “Having” clause.

Order By

In the Order By clause you can provide sorting (ascending/ descending) criteria for the SQL statement. For a report with grouping, the “Order By” clause must have the columns in the same order as the respective sections in the Layout Editor.

**Caution**

An order-by report query that involves millions of events can fail to run and display the following error messages: “The server is too busy, try again later”.

Therefore, HP Recommends that you follow these best practices:

- Use the ‘scan limit’ parameter to limit the number of events that will be scanned.
 - Rewrite the report query to group by name or group by time to reduce the granularity of events scanned.
-

Edit Tab

When you switch from the Design tab to **Edit** tab, the SQL in the Design tab is constructed and displayed as a complete SQL statement in the Edit tab. You can use the Edit tab to view and write more complex SQL statements that cannot be defined in the Design tab.

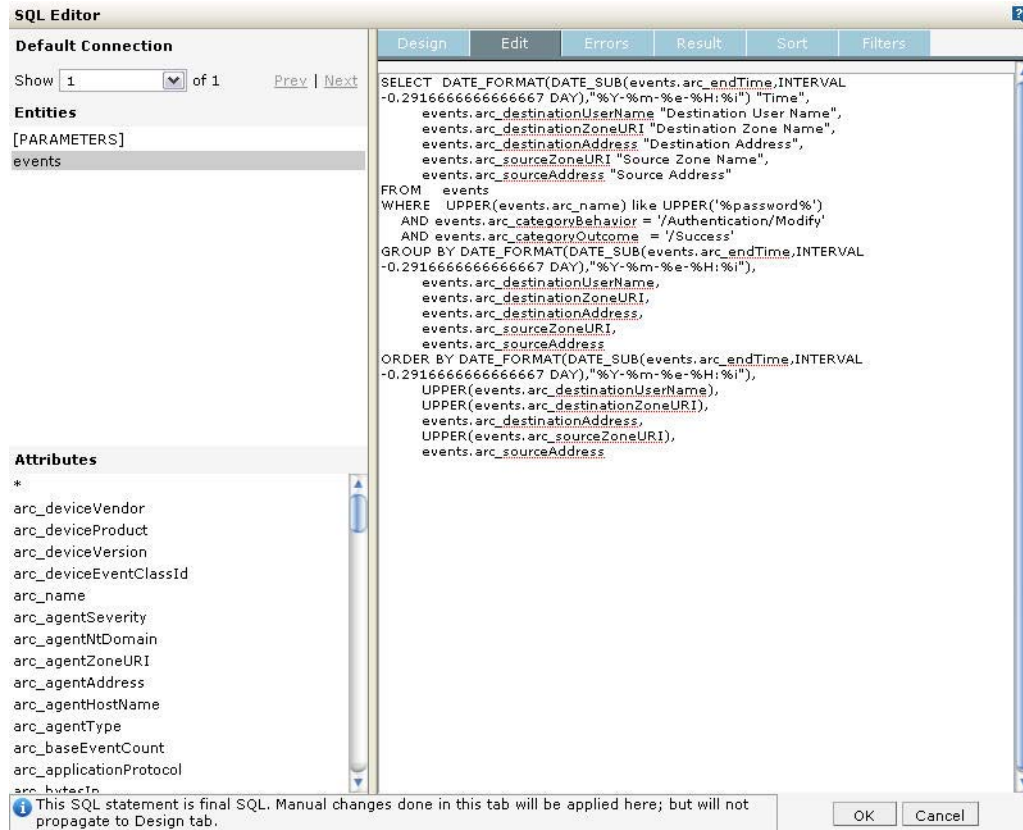


Figure 5-34 SQL Editor: Edit Tab

Relationship of Edit and Design Tabs

The SQL Editor manages the SQL statement being constructed to prevent a complex query (defined in the Edit tab) from being unintentionally overwritten with changes made subsequently on the Design tab.

If you first enter a complex query on the Edit tab, then click back to the Design tab and make changes there, then click the Edit tab again. A dialog prompts to ask whether you want to overwrite the original statement on the Edit tab with the changes you made on the Design tab.

- If you click **OK**, your changes in the Edit tab are overwritten, because the SQL in the Design tab will be reconstructed.
- If you click **Cancel**, the SQL in the Edit tab remains intact and is used as the final SQL. The SQL statement as reflected in the Edit tab will be used as the final SQL for compilation.

Result Tab

The **Result** tab shows query results based on the currently-specified SQL statements (shown in the Edit tab). If the SQL uses a parameter, you will be prompted to provide the values to view the query results.

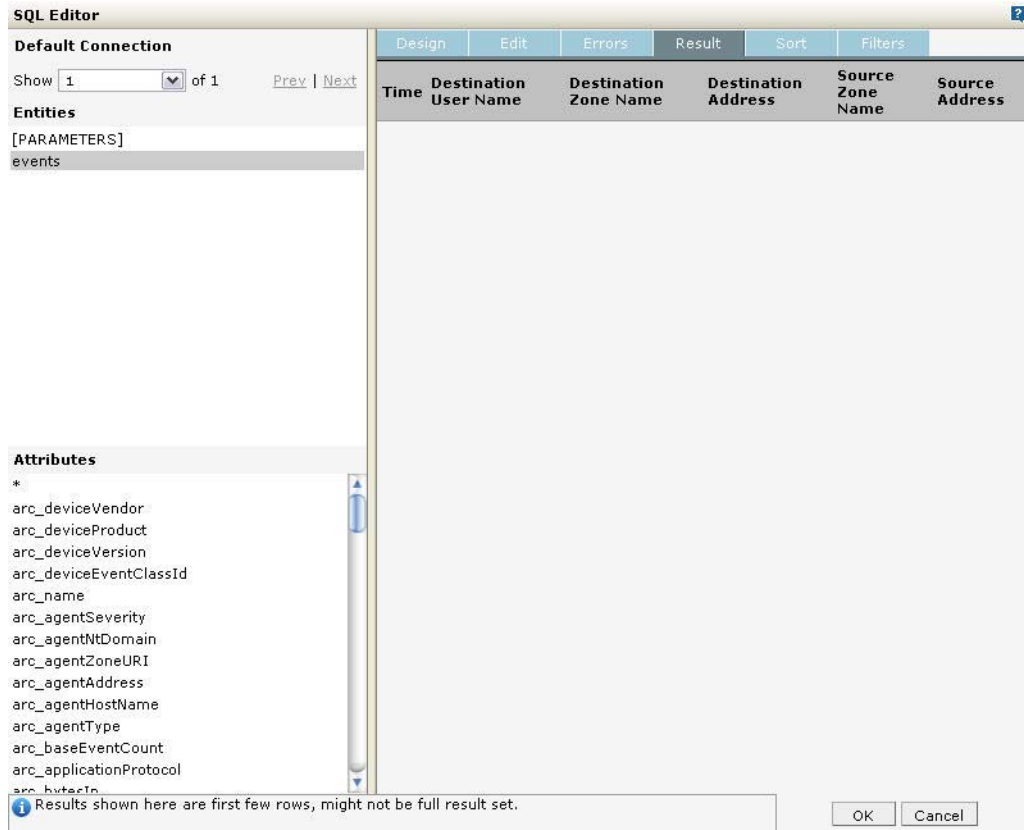


Figure 5-35 SQL Editor: Result Tab

Sort Tab

Click the **Sort** tab to specify levels of sorting at report run time.

SQL Editor Schema refreshed on: 10/12/2011 10:23:00

Entities

[PARAMETERS]

events

Attributes

Design Edit Result **Sort** Filters

Prompt ☐ Check this if you want end user to choose sort fields. The parameter input form shall prompt report fields for user to select.

Count 3 Set the maximum levels of sorting that the end user can set.

Disable Forced Sorting ☐ Check this to avoid sorting for grouping purposes after retrieving from data source. Ensure that the data source (SQL or Procedure) itself returns data in required sorted order.

Select Sort Fields

Available Fields Selected Fields Qualifier

Default Fields

+	Caption	Field	Order
---	---------	-------	-------

Check prompt, if end user should choose sort fields. (Not applicable for Adhoc Reports) OK Cancel

Figure 5-36 SQL Editor: Sort Tab

The following table explains the settings on the Sort Tab.

Table 5-15 Sort Tab Options

Field	Description
Prompt	Check this box if you want the report to prompt for sort order at run time. If Prompt is enabled (checked), at report run time a dialog will pop up to prompt the user to specify a sort order.
Count	Specify the number of levels of sorting you want. For example, if you want to sort by Country, then by State and then by County, select 3.
Disable Forced Sorting	Check this box if you do not want the user to re-order the data once it is sent from the database server.

Filters Tab

Click the **Filters** tab to add filters to a query. This is useful when a report needs to present one or more optional parameters at run time and you want the user or report designer to select the parameter(s) via a multi-select combo box.

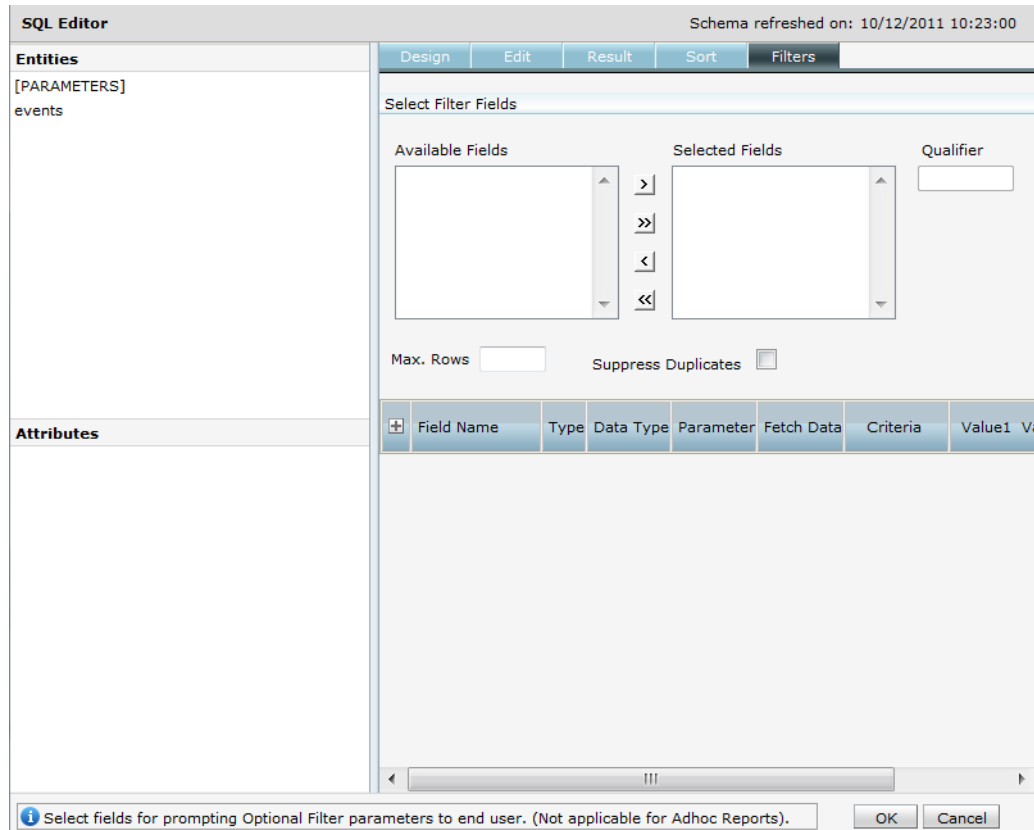




Figure 5-37 SQL Editor: Filters Tab


To get a row at the top:

Click  (Add a filter) in the leftmost cell of column header. This inserts a row at the top.

To get a row below current row:

Click  (Add a filter) in the row below which you want to add a row for condition. A row is inserted below the current row.

To remove a condition:

Click  (Remove this filter) next to a condition you want to delete to remove the filter.

To specify a filter:

Specify field names and associated parameters as described.

Field	Description
Field	Field on which to filter.

Field	Description
Type	Sets the filter type: <ul style="list-style-type: none"> Select UseParameter to determine compare it (equality) with a parameter value that the user specifies at run time. Select ADHOC to allow the user to select condition type at run time.
Data Type	Sets the data type for the parameter: <ul style="list-style-type: none"> CHAR NUMBER DATE
Parameter	In Parameter drop-down box, select the parameter to be used with this filter
Fetch Data	If Fetch Data is selected (checked), the report server will <i>pre-fetch</i> the data, before the parameter form is presented to the user at run time.
Criteria	
Value 1	
Value 2	
Mandatory	

Working with Parameters

Reports get data by running pre-built query objects. If a query needs a value at report run time, it uses built-in, run-time parameters. At report run time, the user is prompted to provide values for run-time parameters as a prerequisite for running the report. The report is then generated based on the user-provided values for those parameters.



We recommend first designing all needed parameter objects before creating the query object that will use those parameter objects. (For information on creating queries, see [“Setting up Queries” on page 201.](#))

Parameters are stored on server and so can be used in one or more report and query objects.

To view and work with Logger Report parameters, click Design | **Parameters** in the Reports left pane or click **Parameter Explorer** and click on a category, select a parameter, and click the Edit Parameter Details button to open the Parameter Object Editor.

Figure 5-38 Report Parameters Object List

Creating New Parameters

To create a new parameter:

- 1 In the Parameter Object Editor, click the **Add New** button located at the top left.
- 2 Specify values for the new parameter. (Details are given in the topics below.)



The parameter name must be unique amongst all parameters in the system.

- 3 After providing all required values, click **Save**.

The parameter is added to the Parameters list.


Setting Parameter Name, Data Type, and Default Values

Specify the parameter unique ID, display name, data type, size, format, and default value as described in the table below.


Table 5-16 Parameter Name, Data Type, and Default Values

Option	Description
Name	Provide a name to uniquely identify this parameter. This name should be unique amongst all parameters in the system.

Table 5-16 Parameter Name, Data Type, and Default Values (Continued)

Option	Description
Prompt	Parameter name displayed on-screen to the user at report run time.
Data Type	Specify type of value the user must provide at report run time: <ul style="list-style-type: none"> CHAR - Value may include alphabetical characters, numbers and special characters. NUMBER - Value may include digits and decimal points DATE - A date or part of a date, like day, month, or year BOOLEAN (For more information, see "To set up a BOOLEAN parameter:" on page 227.)
Size	Specify number of characters or digits this parameter should accept. Note: This is only applicable to CHAR and NUMBER data types, not for Boolean or Date type parameters.
Format	Select the appropriate format in which user should provide value for this parameter. Click  to open a Data Format dialog box. Based on the format you have selected, a format string will appear in the entry box.
Default Value	Specify a default value that is appropriate in most cases to provide for this parameter at report run time. The default value will be automatically selected at report run time. The user can change the default value, if needed. If the user does not change it, the report will run using the default value you specify here for this parameter.

Default Value for Date Type Parameter

For a date type parameter, the **Default Value** field provides a drop-down menu and a calendar. Click the calendar  to provide an explicit date, or select one of these dynamic variable values from the drop-down menu:

- CURRENT_DATE
- MONTH_START_DATE
- YEAR_START_DATE

You can also set a default date that is *relative* to any of the above three dynamic variable dates.

For example, to set a default date as 3 days after CURRENT_DATE, specify
CURRENT_DATE + 3.

To set a default date as 5 days before MONTH_START_DATE, specify MONTH_START_DATE
- 5.

To provide a value that is relative to a dynamic variable, select one of the dynamic variables, then type a suffix to it in the Default Value field by adding + or - and the number.

At report run time, a parameter with a "Date" format will display with the default date set here.

Defining Input Type





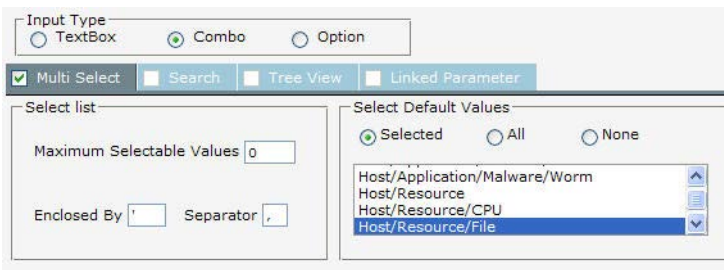
Figure 5-39 Parameter Input Type

The parameter *input type* describes the style of interface provided to users at report run time in which to enter a value for this parameter. Choose from **Text Box**, **Combo**, or **Option** as described below.



In the Reports Designer, changing the parameter type TextBox to another type causes an error. If you need to change the parameter type to TextBox, do not edit an existing parameter, delete that parameter and add a new one.

Table 5-17 Input Type

Option	Description
Text Box	Select “Text Box” input type if you want the user to type the value for the parameter.
Combo	Select “Combo” if you want the user to select one value or multiple values from a drop-down menu. Select the Multi Select checkbox so that user can select multiple values from the box. 
Option	Select “Option” if you want the user to select values represented as options. Select the Multi Select checkbox to have value options in the form of checkboxes. Keep Multi Select checkbox clear to have options in the form of radio buttons.

Setting up Boolean Parameters

Parameters that have a Boolean “Data Type” are represented to the user as checkboxes (the input type) and have only two states:

- Checked (chosen at run time)
- Unchecked (de-selected at run time)

To set up a BOOLEAN parameter:

- 1** Select **Data Type** as BOOLEAN.
- 2** In the **Values** area, for **Checked** specify the value to be passed when the user chooses this option at run time (selects/checks the checkbox presented).
- 3** In **Unchecked** specify value to be passed when the user does not choose this option at run time (de-selects/leaves the checkbox unchecked).

Setting Various Run Time Behaviors

You can specify a variety of options on how the parameter will look and act at report run time. These options are generally related to the input type, and further define acceptable user input values, whether the parameter will be displayed or hidden, provide searchable values, and so forth.

Table 5-18 Parameter Options

Option	Description
Mandatory	Select this checkbox if you want to require the user to specify a value for this parameter at report run time.
Visible	Select this checkbox if you want the parameter to be visible (displayed) on the input form at report run time. Keep this unchecked (clear) if the value for this parameter is populated from another report or if you want the parameter to use the default value in all cases.
Restrict to List	This setting is applicable for parameters with Input Type of Combo . Select (check) the Restrict to List checkbox here to force user input of a parameter value from the available run-time options only. If Restrict to List is <i>not selected</i> in the parameter definition you create here, the user can specify a value or can select value(s) from available options.
Pass Values Using Tables	This setting is applicable for Multi Select . Select this checkbox when you want to pass parameter values through a table. This is done especially when the number of values that can be passed (total number of bytes of selected values) as part of the SQL is more than allowed.
Enable	
Forced	Select this checkbox if you want to restrict parameter values to a pre-specified list of values.

Setting the Data Source List

Specify values for Checkbox, Combo, and Option input type. Values can be predefined only.

- *Separator*—Specify the character to use to separate the two values. This will depend on the database.
- *Select Default Values*—Specify the number of default values to display at report run time. You can choose from
 - ◆ Selected—Only values for the selected parameters are displayed.
 - ◆ All—Values for all parameters are displayed.
 - ◆ None—No values are displayed. That is, no default values are defined.

Modifying a Parameter

To modify a parameter:

- 1 On the **Reports** right panel menu, click **Parameter Explorer** to bring up the Parameter Object list.
- 2 In the **Parameter Objects** list, select the name of parameter that you want to modify and click the **Edit Parameter Details** button.
- 3 Edit the parameter as needed (via the settings described in [“Creating New Parameters” on page 224](#)) and click **Save**.

Only custom parameters can be modified, not supplied parameters, since supplied parameters are required for use in system Reports and Solution pack add-ons.

Deleting a Parameter

To delete a parameter:

- 1 On the **Reports** left panel, click **Parameter Explorer** to bring up the Parameters Object list.
- 2 In the **Parameters** list, select the name of parameter that you want to delete.
- 3 Click **Delete**.

Only custom parameters can be removed, not supplied parameters, since supplied parameters are required for use in foundation Reports and Solution pack add-ons.

Configuring Parameter Value Groups

Some reports may require users to provide multiple run-time values that would be easier to select if they were grouped. For example, a report that requires a user to select more than one country name might be a good candidate for parameter value groups. Users might find it difficult to select a few country names from a single, long list of countries.

As an alternative, the query designer could create parameter value groups for the Americas, Europe, Asia, Africa, and so forth; each with lists of countries belonging to those continents or areas. At report run time, when the user selects a group, values belonging to the group are pre-selected. Users do not have to manually select countries in, for example, Europe or Asia, for every report run. Selections are saved from one report run to the next.

Using parameter value groups as a part of your query design strategy can save users time and reduce error at report run time.

To view and work with Logger Report parameter value groups, click Design | **Parameter Value Groups** on the Reports left panel.

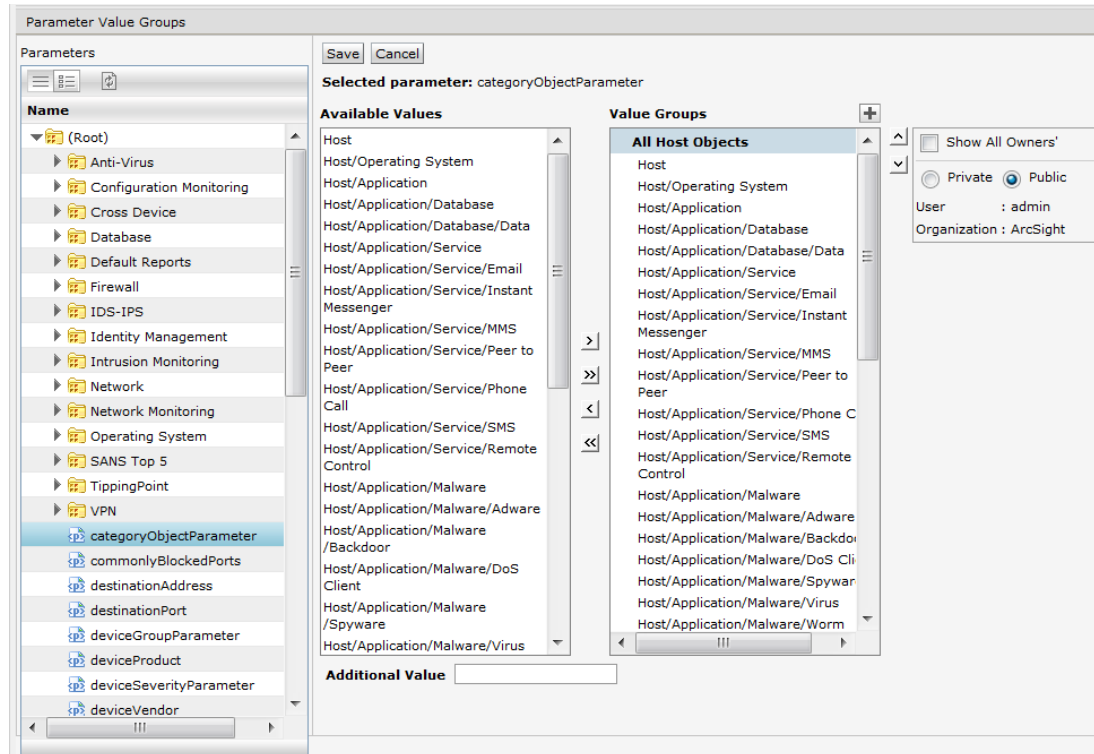



Figure 5-41 Parameter Value Groups

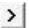
The following table describes the options on the Parameter Value Groups page.

Table 5-19 Parameter Value Groups



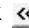
Option	Description
Parameters	Lists all the parameter objects.
Available Values	Lists available values for the selected parameter.
Value Groups	Lists groups created and the values selected within a group. An icon appears on the left of a Private group.
Show All Owners	If selected, displays groups created by all users.
Option buttons: Private Public	Select Private to list the groups you have set for you only. Select Public if you wish to list the groups you have set for everyone.





To create a group:

- 1 Click  (Add Group) next to the **Value Groups** box. A group is created and listed under Value Groups with a default name (based on the currently selected parameter in Parameters list).
- 2 In the Value Groups list, edit the new group name as needed. (Double-click the name to edit it, if it is not already in edit mode.) Double-click the name again to set it, or click outside the box.

- 3 Add the values you want in the group by selecting a value in **Available Values** list and clicking  (Add value to selected group) button. The selected value is added to the selected group in the Value Groups list.
- 4 Repeat [Step 3](#) for each value you want to add to the group.

If a value that you want to add to a group is not listed in Available Values list, specify the value in **Additional Value** field (under Available Values) press Return key. The custom value is added to the currently selected group.

Select an Available Value and click  to add all the values to the selected group in Value Groups, click  to remove the selected value from Value Groups, and click  to remove all the values from Value Groups box.


Select a group and click up  and down  arrows to move the selected group up or down. Select a value and click up  and down  arrows to move the selected value up or down (within the group).

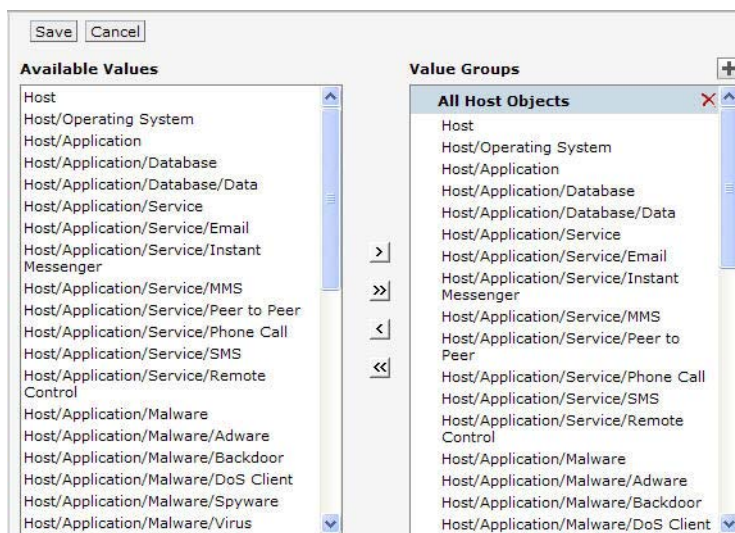
- 5 Click **Save**.




If the name of a group is changed by a user, the values under that group will be removed from the "Selected Values" group of that user's preferences.


To create a tree view parameter:

To select a value, click the leaf node and click  button.



To select all values in a branch (only for a multi-select parameter), click the respective branch and click the  button. All the values under that branch will be selected.

To make changes in name of a group, double-click the group name to make it editable. Specify a new name and click outside the box.

To delete a group, click  in the title of group you want to delete and click the **Save** button to save the changes.

Applying Report Template Styles

Logger Reports use a style file (.sty) to generate report output per a specified format. The style file defines the look and feel, arrangement, orientation, and so on, of the report output.

You can modify any of the style files from the Logger Reports Template Styles page or you can define a new style to suit your needs.



A report layout file (.irl) defines factors like paper size, static controls, headers and footers to include in a report, and so on. You can define your own layout files. See [“Defining a New Template” on page 232](#) for more information.

To view and work with Logger Report template styles, click Design | **Template Styles** on the Reports left menu bar.

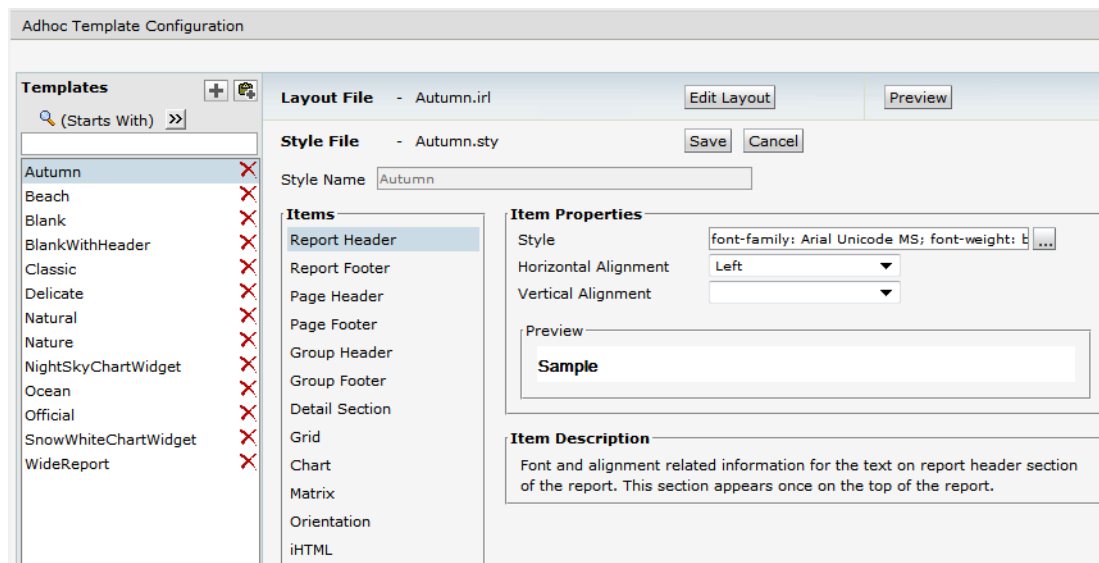


Figure 5-42 Report Template Styles Configuration


Defining a New Template

Before creating a new template, you may want to check whether there is an existing one that meets your needs.

To search for an existing template, do one of the following:

- Enter the first few letters with which the template name begins (if the “Starts With” search criteria is selected) in the text box above the list of existing templates.
- Enter a word or part of a word that the template name contains (if the “Contains” search criteria is selected) in the text box above the list of existing templates.

To define a new template:

- 1 Click Design | **Template Styles** on the Reports left menu bar.
- 2 Click the  icon in the right panel.
- 3 Define the Items and Item Properties for the template.

- 4 If you want to define or change the report layout file, click **Edit Layout**.



You will need to edit the layout of the report to include a header or footer in a report. After clicking Edit Layout, click "Report Header" (to include a header) or "Page Footer" (to include a footer) to select that section. Click **Insert > Layout Control > select an option from the sub-menu**.

- 5 Click **Save**.

Scheduling Reports

You can schedule reports to run as scheduled "jobs" on a one-time basis in the future, or set a frequency schedule (hourly, daily, and weekly). As part of scheduling a report job, you can set delivery options to e-mail, save, or publish the resulting reports.

HP recommends using the Scheduled Report feature in lieu of running on-demand (ad hoc) reports whenever possible, so that reports are run during periods of light load. For more on this see ["Best Practices" on page 172](#).

Make sure you are familiar with the information in ["Impact of Daylight Savings Time Change on Logger Operations" on page 363](#) before you schedule reports.

Viewing and Editing Scheduled Reports

Reports that are already scheduled are displayed on the Scheduled Reports page.

To view scheduled reports:

Click **Scheduled Reports** on the Reports page left pane to view a list of currently scheduled jobs.




To view scheduled reports, a user must belong to a Logger Reports Group, a Logger Search Group, and a Logger Rights Group.

Add				
Task	Type	Schedule	Next Run Time	
Password Changes	Report	Sunday at 23:00	Sun Sep 30 23:00:00 PDT 2007	 
Top 10 Talkers	Report	Saturday at 23:00	Sat Sep 29 23:00:00 PDT 2007	 
Top User Logins	Report	Daily at 23:00	Sat Sep 29 23:00:00 PDT 2007	 

Figure 5-43 Scheduled Reports

To edit a scheduled report:

Click  (Edit) next to the scheduled report job you want to edit.

This brings up the Edit Report Job page, which lets you change most of the settings on the scheduled job. Modify the settings as needed and click **Save**.

For details on how to specify these settings, see [“Scheduling a Report” on page 234](#).




Note

The job name is not editable once the scheduled report job is created.

Other settings can be modified with an edit, and work the same way as on the Add a Report Job page described in [“Scheduling a Report” on page 234](#).

To remove a scheduled report:

Click  (Delete) next to the scheduled report job you want to remove.



Tip

Removing the report from Scheduled Reports list here deletes the *scheduled job*, not the report itself nor any instances of its previously published output.

Scheduling a Report

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 363](#) before scheduling a report.

To schedule a report:

- 1 Click **Scheduled Reports** on the Reports page left menu.

The page shows the list of currently scheduled report jobs, if any. (See [Figure 5-43](#).)

- 2 Click **Add** to bring up the Add Report Job page.

- 3 On the Add Report Job page, use the drop-down menu next to **Report Name** to select a report, and click **Go** to load the report.



Note

You must click **Go** to load the selected report at the Report Name field before you save the scheduled report job. Attempting to save the scheduled job without first loading the report name will result in an error, and the report will not be saved.

- 4 Choose one or both delivery options under Delivery Option (**Email**, **Publish**). By default, only Publish is selected.

To keep the Publish option, enter its associated parameters. Uncheck the option to unselect it. If you also want the Email option, click it and enter its associated parameters. Both E-mail and Publish options for scheduled reports are the same as those provided after you run a report “on demand”.

- ◆ **Email** - For details on setting e-mail delivery options, see [“E-mailing a Report” on page 182.](#)

The screenshot shows the 'Delivery Operations' form with the 'Email' tab selected. The 'Report Name' field is set to '(Select Report)'. Under 'Send Report As', 'Link' is selected. The 'Save In' field is empty, and 'User's Working Folder' is checked. The 'File Name' field is empty, and 'Suffix Timestamp Format' is unchecked. The 'To' field is empty. The 'Cc' field is empty, and the 'Bcc' field is empty. The 'Subject' field is empty. The 'Message' field contains the text: 'Hello,

 You have received this automated email to let you know that report <%MENU_NAME%> has been generated.Please click the following link to view the report in'. On the right, the 'Report Format' is set to 'ACROBAT PDF'. Under 'Delivery Options', 'Horizontal Breaks' is selected, and 'Deliver Zipped' is checked.

- ◆ **Publish** - For details on setting publishing options, see [“Publishing Reports” on page 180.](#)

The screenshot shows the 'Delivery Operations' form with the 'Publish' tab selected. The 'Report Format' is set to 'ACROBAT PDF'. Under 'Delivery Options', 'Horizontal Breaks' is selected, and 'Deliver Zipped' is checked. The 'Save In' field is empty, and 'User's Working Folder' is checked. The 'File Name' field is empty, and 'Suffix Timestamp Format' is unchecked. The 'Valid Upto' section has three options: 'Public' (selected), 'Private' (unchecked), and 'Date' (unchecked). The 'Valid Upto' section also has a 'Valid Upto' field set to '1' month, and a 'Valid Upto' field set to '2/22/2014'.

- 5 Fill in the rest of the fields based on the report you chose, as described in [“Add Report Job Settings” on page 236.](#)
- 6 Click **Save**.

The report you added is scheduled, and now shows on the Scheduled Reports list.



If you got a batch error when you clicked Save, try clicking Go next to the Report Name to reload the report per [Step 3](#). This is the most common oversight in terms of specifying the job parameters.

Add Report Job Settings

The following table describes the Add Report Job settings.

Table 5-20 Add Report Job Settings

Option	Description
Name	Provide a name for the report job. This is the name that will be displayed on the Scheduled Jobs list.
Schedule	<p>Set the frequency for the scheduled run of the report.</p> <p>For example, you can specify to run the report on specified "Days of the Week" like Sa, Su, M, T, and so forth, or "Everyday".</p> <p>You can choose to run the report at a certain hour every day "Hour of the Day" or "Every" hour so many hours.</p>
Report Name	<p>Select a report from the list, and click Go to load the report.</p> <p>Note: You must click "Go" to load the selected report at the Report Name field before you save the scheduled report job. Attempting to save the scheduled job without first loading the report name will result in an error, and the report will not be saved.</p>
Delivery Options	<p>Depending on which delivery option you choose, the associated parameters are displayed. Click to enable (check) or disable (uncheck) these options.</p> <p>Both E-mail and Publish options for scheduled reports are the same as those provided after you run a report "on demand".</p> <p>Select a delivery option:</p> <ul style="list-style-type: none"> • Email - For details on setting e-mail delivery options, see "E-mailing a Report" on page 182. • Publish - For details on setting publishing options, see "Publishing Reports" on page 180. • Save In • File Name • Suffix Timestamp Format • Valid Upto • Pagination
Report Format	<p>Select a report format (Acrobat PDF, HTML, and so forth.)</p> <p>For details on report formats, see "Report File Formats" on page 179.</p>
Report Parameters	<p>You can either accept the default parameters, or modify them here. These are the same parameters that can be specified for an on-demand report run.</p> <p>For information on specifying report parameters, see "Quick Run with Default Options / Run In Background Report Parameters" on page 175.</p>

Deploying a Report Bundle

You might obtain additional sets of reports from ArcSight to address new security scenarios, add packaged solutions, or enhance your current coverage with updated reports. You can use the **Deploy Report Bundle** page to load and deploy packages of new reports onto your Logger system.

On the Reports page left panel menu, click the **Deploy Repository Bundle** link to get started.

Figure 5-44 Deploy Report Package

A report package (or “cab” file) can contain several types of reporting resources, including:

- Categories and reports
- Organization information
- Schedules
- Portal properties and server properties
- Parameter objects
- Query objects
- Adhoc report templates
- Printer settings
- Database connections

To upload and deploy report package:

- 1 In the entry box provided under Step 1, specify the reports package file name and with its full path. Click the **Browse** button to locate the file.
- 2 Click **Upload**.

The content is uploaded and information is displayed about the included categories and reports. (A legend is provided below these steps).
- 3 If you want to create log of the deployment process, click (check) the **Create Log File** option.
- 4 Click **Deploy** to continue with the deployment process, or click **Cancel** to discontinue with deployment process.)

Status information is displayed about the objects in the package being deployed.

A legend is displayed just below the Deploy button. Information about each of the components in the package is displayed in respective tabs.



Note

Overwrite behaviors are determined when the package was created.

For example, protocol on whether or not an object in the deployed package will overwrite an existing object on the system, and under what circumstances, is determined at package creation time. Therefore, these settings on package deployment are not available to you at deploy time.

A log file will be created if the “Create Log File” checkbox was selected.

The content of the deployed reports package is available on the respective Logger Reports pages. Solution Reports will be listed under **Solution Reports** on the left panel menu. For more information about these types of reports, see [“Solution Reports” on page 150](#).

Report Server Administration

Logger Reporting provides a default configuration for the report server. If you do not modify the report server, reports will run with the default settings.

Timeouts when Running Reports

There are two timeouts that can affect long running reports. The client timeout is 1 hour. If an Adhoc report takes more than an hour to run, it will time out. Use a scheduled report instead. The default database connection timeout for scheduled reports is 4 hours. If a scheduled report takes more than 4 hours, to run, you can increase the database connection timeout from the Report Configuration pane.

Report Configuration

To view or modify the report server configuration:

- 1** Click **Reports** from the top-level menu bar.
- 2** Click **Report Administration** in left panel menu. The Report Configuration pane opens.

Report Configuration	
<input type="button" value="Save"/> <input type="button" value="Cancel"/>	
Database Connection TimeOut (seconds)	<input type="text" value="14400"/>
Log Level	ERROR ▼
Data Source Fetch Size (rows per fetch)	<input type="text" value="50"/>
E-Mail From Address	<input type="text"/>
Host URL	<input type="text" value="https://<logger_hostname>/logg"/>
SMTP Server	<input type="text" value="127.0.0.1"/>
<input type="button" value="Save"/> <input type="button" value="Cancel"/>	

The following table describes the report configuration settings.

Table 5-21 Reports Configuration

Option	Description
Database Connection Timeout	<p>Time in seconds after which the database connection will be closed, if not used for that many seconds.</p> <p>Valid values for this timeout include any integer greater than zero.</p> <p>Default: 14400</p> <p>Example: If DATABASE_CONNECTION_TIMEOUT is set to 50, report server will close the connection to a database if there is no communication between the report server and database server for 50 seconds.</p>
Log Level	<p>Sets the level of criticality to be considered for logging.</p> <p>Valid values are DEBUG, INFO, WARN, ERROR, FATAL.</p> <p>Default: ERROR.</p> <p>Example: LOG_LEVEL = ERROR</p>

Table 5-21 Reports Configuration (Continued)

Option	Description
Data Source Fetch Size	<p>Specifies the number of records to be fetched from the data source at one time (in one “read”).</p> <p>A valid value is any positive integer.</p> <p>Default: 50</p> <p>Example: DATA_SOURCE_FETCH_SIZE=50</p>
E-mail from Address	<p>Sets the e-mail address to be displayed as the “from” (sender’s) address in e-mails originating from the Logger Reporting system.</p> <p>Default: None.</p> <p>Example: loggeradmin@companyxyz.com</p>
Host URL	<p>Host URL (URL to be specified to run the Logger application) sent as part of Logger Reporting e-mails.</p> <p>Syntax: HOST_URL=[Host URL](String)</p> <p>Default: https://<logger_hostname>/logger/report</p> <p>Example: HOST_URL=https://loggerA.companyxyz.com/logger/report</p>
SMTP Server	<p>Sets the server IP address or domain name (as IP or URL) used to e-mail scheduled reports. All e-mail communications, such as notifications and report delivery, are sent by Logger Reporting via this e-mail server.</p> <p>Example: SMTP_SERVER=127.0.0.1</p> <p>For information on Logger’s SMTP settings, see “SMTP” on page 363 for Logger appliances and “SMTP” on page 414 for software Loggers.</p>

Using Report Category Filters

A Search Group filter can be optionally assigned to each report category. Assigning a Search Group filter to a report category means that all the reports in the category will only process events returned by this filter.

To assign a search group filter to a report category:

- 1 Create the filter that you would like to apply to every report in a given category. See [“Filters Tab” on page 312](#) for the details of creating a filter of type Search Group.
- 2 Click the **Reports** tab. In the menu, under Administration, click **Report Category Filters**.
- 3 The new search group filter will appear in the pull-down menu associated with each category. Select the desired filter for each category.
- 4 Click **Save**.

To remove a search group filter from a report category:

- 1 Click the **Reports** tab. In the menu, under Administration, click **Report Category Filters**.
- 2 In the pull-down menu associated with the report category from which you want to remove the filter, select **None**.
- 3 Click **Save**.

Backup and Restore of Report Content

You can backup and restore report content. For more information about this feature, see [“Configuration Backup and Restore” on page 330](#).

iPackager

The iPackager is a utility that allows you to package all or selected reports and/or report objects residing in Logger. This package can be later imported on a different Logger installation at any time. If you own multiple Loggers, you can use the packages created by iPackager to configure the reporting feature on them. Set it up at one location and create a package using the iPackager, then you can upload the package at other locations. This method completely eliminates the need of re-doing all the configuration activities at multiple locations.



Note

- The iPackager utility can only be used by users with administrator privileges.
- Java must be enabled on your browser for iPackager to run.

To access the iPackager, click the **iPackager** link in the left panel of the Reports page.

Using the iPackager, you can package and distribute:

- Repository objects like folders, reports, query objects, parameter objects, OLAP layouts, and dashboards (along with access right information)
- Organization, user, user access rights, and user mappings
- Schedules
- Data connection information (along with access rights information)
- Print settings
- Web client properties
- Report server configuration files
- Adhoc report templates
- Java plug-ins
- Approval process details

The iPackager allows you to first create a .conf file in which you can collect (import) the references for all the components that you want to include in the package. You can save the .conf file and edit it at any time. Once you are satisfied with the contents of the .conf file, you can build the package. The package is built into a CAB file. The components referenced in the .conf file are actually picked up when building the CAB. Data to be packaged in a single CAB can be imported from multiple report servers.

You can open only one .conf file in iPackager at a time. When iPackager opens a .conf file, it checks for the availability of the components already imported in the .conf file. If any of the components already imported are not found on the report server, it is indicated on the tree-view. In such a case, a CAB cannot be built.





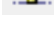


The iPackager Page

The iPackager page consists of three panes that are located to the right of the navigation pane on the Reports page. The pane to the immediate right of the navigation pane in the

reports page displays the presently included contents of a .conf file that is open in the iPackager. These contents are displayed in a tree view. The **Build Properties** box displays the details of the object selected from the CAB. You can edit what contents you would like to include in the CAB file in the Build Properties box using the buttons described below. The Messages and Problems tabs display messages, logging details and problems faced while creating or editing a package.

Buttons Available from the iPackager

The following buttons available on top of the iPackager page represent the various actions that can be performed from within the iPackager:

Buttons	Description
	New Creates a .conf file.
	Open Opens an existing .conf file in iPackager.
	Save Saves the currently open .conf file.
	Save As Saves the .conf file that is currently open under a new name.
	Clear Package Clears the contents of a package that is currently open in iPackager.
	Import Selective Data from Report Server Imports data as references into the .conf file that is currently open in iPackager. You can specify what you would like to import. For example, you can choose to import query objects parameter objects and reports only.
	Import Complete Data from Report Server. Imports everything from the report server.
	Cancel Server Import Cancels a current import action from the report server.
	Build CAB Initiates the process of building a CAB file.
	About Gives you the version number of the iPackager

When you first open the iPackager page or when you click the New button, you will be prompted to enter the following identifying information for the CAB file that you will be creating with the currently displayed .conf file:

- Author
- Company

- Version
- Comment

Importing References from the Report Server

You can import references from a report server into a .conf file. Keep in mind that only references to those components will be imported into the currently open .conf file in the iPackager. The actual components will be picked up during the creation of the CAB.

To import selective data:

- 1 Click the **Import Selective Data from the Report Server** button located on top of the iPackager page. You will see the following form:

The screenshot shows a web form with two main sections. The first section, titled 'Repository', contains a grid of checkboxes for selecting data to import: 'Categories', 'Reports', 'Parameter Objects', 'Query Objects', 'Include Value Groups', 'Dashboards', and 'Dashboard Widgets'. The second section, titled 'Configuration', contains a single checkbox for 'Templates'.

- 2 Check whatever components you want to import.
- 3 You will be prompted to select specific pieces within the component for example, if you checked the Reports checkbox, you will further be prompted to select the specific reports you want included.
- 4 Once iPackager completes importing the selected data, you can see the references to that data in a tree structure in the left pane of the iPackager page.
- 5 Click **Save** or **Save As** icons to save the .conf file.



Note

Since Reports are saved in categories, if you select only a report, its category gets automatically selected too.

Click the **Cancel Server Import** button from the toolbar at any point of time to cancel the data transfer. In this case, the .conf file state will be restored to the state before the import began. None of the data from the current import will be included in the package.

To import everything from the Report server:

- 1 Click the **Import Complete Data from Report Server** button. You will see the same page as shown above with every checkbox checked for you.
- 2 You still have the option to uncheck any box against any component if you decide you do not want to import it.
- 3 Click **Import**.
- 4 Once iPackager completes importing the selected data, you can see the references to that data in a tree structure in the left pane of the iPackager page.
- 5 Click **Save** or **Save As** icons to save the .conf file.

You can see the status of the import in the **Messages** tab located below the Build Properties box.

Modifying Properties for Imported Objects

You can modify component details as well as delete one or more components in an open conf file. To do so, click on the object in the navigation tree in the iPackager page. The properties page for the object appears in the upper right pane. All of the following properties allow you to choose any of the following deployment action on the target chosen:

- **Replace if present** - While importing, if the component is found in the package, replace the one in package with the one on the report server.
- **Add if not present** - While importing, if the component is not found in the package, add it to the package.
- **Delete if present** - While importing, if the component is found in the package, delete it.
- **Cascade Delete** - For Categories only, delete the folder (category) even if it contains reports.

Category Properties

When you click on a Category in the navigation tree of the iPackager page, the following property page opens.

Parent Category

Category

Deployment action on target repository

<input checked="" type="checkbox"/> Replace if present	<input type="checkbox"/> Delete if present
<input checked="" type="checkbox"/> Add if not present	<input type="checkbox"/> Cascade Delete

The **Category** box come pre-filled with the category name that is present on the report server. You can change the name of the category. If you change the name here, the category is packaged with the new name, but its original name on the report server will not change.

Report Properties

When you click on a report in the navigation tree of the iPackager page, the following property page opens.

Category Name	Cross Device	
Report	XD-Config-Configuration Changes by Type	
Path	//127.0.0.1/45450/ArcSight/admin/6172637369676874	<input type="button" value="Browse"/>
Version	0.00	
Deployment Type	CUSTOM	
Deployment action on target repository		
<input checked="" type="checkbox"/> Replace if present	<input type="checkbox"/> Delete if present	
<input checked="" type="checkbox"/> Add if not present		
<input type="button" value="Update"/>		<input type="button" value="Delete"/>

The **Report** box come pre-filled with the report name that is present on the report server. You can change the name of the report. If you change the name here, the report is packaged with the new name, but its original name on the report server will not change.

Query Properties

When you click on a query in the navigation tree of the iPackager page, the following property page opens.

Category Name	
Query Object	DM-Cross-Configuration Changes by User
Deployment action on target repository	
<input checked="" type="checkbox"/> Replace if present	<input type="checkbox"/> Delete if present
<input checked="" type="checkbox"/> Add if not present	
<input type="button" value="Update"/>	<input type="button" value="Delete"/>

The **Query Object** box come pre-filled with the query object name that is present on the report server. You can change the name of the query object. If you change the name here, the query object is packaged with the new name, but its original name on the report server will not change.

Parameter Properties

When you click on a parameter in the navigation tree of the iPackager page, the following property page opens.

Category Name	
Parameter Object	tpFilterCnt
Deployment action on target repository	
<input checked="" type="checkbox"/> Replace if present	<input type="checkbox"/> Delete if present
<input checked="" type="checkbox"/> Add if not present	
<input type="button" value="Update"/>	<input type="button" value="Delete"/>

The **Parameter Object** box come pre-filled with the parameter object name that is present on the report server. You can change the name of the parameter object. If you

change the name here, the parameter object is packaged with the new name, but its original name on the report server will not change.

Template Properties

When you click on a template in the navigation tree of the iPackager page, the following property page opens.

File Name:

Deployment action on target repository

☒ Replace if present ☐ Delete if present

☒ Add if not present ☐ Delete if not present

Opening a .conf File

To open an existing file in iPackager:

- 1 Click the **Open** button in the toolbar on top of the iPackager page.
- 2 Navigate to the saved .conf file and click **Open**.

Deleting an Item from the .conf File

To delete an item:

- 1 Open the .conf file in iPackager.
- 2 From the iPackager page, navigate to the item and select it. Selected item's details will be displayed on the right pane (properties area).
- 3 Click the **Delete** button in the Properties area. A warning dialog will appear.
- 4 Click **Yes** to confirm deletion of the selected item.

Clearing the Contents in a .conf File

To clear the contents:

- 1 Open the .conf file in iPackager.
- 2 Click the **Clear package** button in the toolbar.
- 3 A warning dialog box appears.
- 4 Click **Yes** to go ahead with deletion.

Building the CAB

When you issue command to build the cab, the actual objects specified in the references in your open .conf file are actually picked up from the respective locations and a CAB file is built. This CAB file will contain all the objects.

If any of the information saved in the .conf file is not available at the right source while building the CAB, then you will see an error message and the CAB building process will be stopped. You will need to fix any errors before rebuilding the CAB file.

To build the cab:

- 1 Click the **Build CAB** button on the toolbar. The Build Properties screen appears.
- 2 Specify the build properties and click the **Build Cab** button. The Save dialog box opens.
- 3 Specify the credentials and location to save the file.
- 4 The Cab building process will begin.

Deploying a CAB file in Logger

**Caution**

When deploying a CAB file from a source Logger to a target Logger, if the categories being imported do not have identical names and IDs on both Loggers the deployment will fail.

Should you run into this issue, rename the conflicting category in the target Logger or the source Logger (you will need to recreate the CAB file if you do this on the source Logger) such that the category has a unique name and/or ID. Then, redeploy the CAB file.

To deploy a CAB file in Logger:

- 1 Click the **Deploy Report Bundle** link in the left pane of the Reports home page. You will see the following page.

- 2 Click the **Browse** button to select the CAB file to be uploaded, and click **Upload**.
- 3 Click the **Deploy** button.

Chapter 6

Configuration

This chapter describes how to create and manage receivers, forwarders, devices, device groups, and filters. Receivers, devices, and other resources created by one user are visible to all other users, although subject to user group privileges. Resources are shared by all sessions.

This chapter includes information on the following areas of Logger configuration:

- ["Devices" on page 247](#)
- ["Event Archives" on page 250](#)
- ["Storage" on page 256](#)
- ["Event Input" on page 260](#)
- ["Event Output" on page 284](#)
- ["Alerts" on page 296](#)
- ["Scheduled Tasks" on page 309](#)
- ["Filters" on page 312](#)
- ["Saved Searches" on page 315](#)
- ["Search" on page 320](#)
- ["Peer Loggers" on page 326](#)
- ["Configuration Backup and Restore" on page 330](#)
- ["System Maintenance" on page 333](#)
- ["Retrieve Logs" on page 352](#)
- ["Content Management" on page 353](#)

Devices

The Devices section manages both Devices and named collections of devices called device groups.

Devices

A device is a named event source, comprising of an IP address (or hostname) and a receiver name. Two receivers can receive events from the same IP address, so IP address alone is insufficient to identify a device. Event source is the device that directly sends the event to Logger. When an event is sent through a SmartConnector, the event source is the system on which the SmartConnector is running and not the device that sent the event to the SmartConnector.

Devices can be added to device groups, and device groups can be referenced in filters and queries. Receivers perform *autodiscovery* by automatically creating a device for each source IP address. Devices created by autodiscovery are named for their hostname, or if the hostname cannot be determined, their IP address.

Figure 6-1 shows the Devices page, which displays all defined devices and includes controls to add, edit, or delete them.

Devices						
Device Groups						
Add						
Name	IP Address	Receiver	Creator	Last Editor		
15.214.128.134 [smart]	15.214.128.134	smart	System			
15.214.131.117 [TCP Receiver]	15.214.131.117	TCP Receiver	System			
15.214.147.123 [TCP Receiver]	15.214.147.123	TCP Receiver	System			
15.214.147.123 [UDP Receiver]	15.214.147.123	UDP Receiver	System			
15.214.157.89 [smart]	15.214.157.89	smart	System			
Logger Internal Event Device	127.0.0.1	Not Applicable	System	System		
Logger Internal Event Device [Apache URL Access Error Log]	127.0.0.1	Apache URL Access Error Log	System			
Logger Internal Event Device [Var Log Messages]	127.0.0.1	Var Log Messages	System			

Figure 6-1 Devices page

Maximum number of devices that can be defined on Logger: No limit.


To pre-define a device:

Autodiscovery creates devices automatically, but you can also pre-define them manually.

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Devices** in the sub-menu, and then click the **Devices** tab. A display similar to that shown in Figure 6-1 appears.
- 3 Click **Add**.
- 4 Enter a name, an IP address, and select a receiver for the new device.
- 5 Click **Save** to add the new device, or **Cancel** to abandon it.

One reason for editing a device is to replace the default name created by autodiscovery (the IP address or hostname) with a more meaningful one.

To edit a device:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Devices** in the sub-menu, and then click the **Devices** tab. A display similar to that shown in Figure 6-1 appears.
- 3 Locate the device that you want to edit and click the Edit icon () on that row.
- 4 Change the Name or IP address for the device.
- 5 Click **Save** to update the device group, or **Cancel** to abandon your changes.

To delete a device:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Devices** in the sub-menu, and then click the **Devices** tab. A display similar to that shown in Figure 6-1 appears.

- 3 Locate the device that you want to delete and click the delete icon (✖) on that row.
Deleting a device does not block the source IP address from sending events. If new events are received, autodiscovery recreates the device.
- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the device.

Device Groups

Device groups allow you to categorize named source IP addresses called devices. The Device Groups page lists all device groups with edit and delete icons and includes the ability to create new device groups.



Device groups can be associated with storage rules that define in which storage group events from a specific devices are stored. Doing so enables you to retain event data from different sources for different lengths of times (because you can define different retention policies on different storage groups). For more information about storage rules, see [“Storage Rules” on page 259](#).

Maximum number of device groups that can be created on Logger: No limit.

To create a device group:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Devices** in the sub-menu, and then click the **Device Groups** tab.
- 3 Click **Add**. A display similar to that shown below appears.

Add Device Group

You may assign one or more devices to a device group.

If you wish to add a device which is not yet created, you must first go to the [Devices](#) page and create it.

To select or deselect devices, ctrl-click each device name.

Name


Devices

- 15.214.128.134 [smart]
- 15.214.131.117 [TCP Receiver]
- 15.214.147.123 [TCP Receiver]
- 15.214.147.123 [UDP Receiver]
- 15.214.157.89 [smart]
- Logger Internal Event Device [Apache URL Access Error Log]
- Logger Internal Event Device [Var Log Messages]


Use ctrl-click to select or deselect items

- 4 Enter a name for the new device group. Click to select devices from the list. Press and hold the **Ctrl** key when clicking to add additional devices to the selection. To select a range of devices, click to select the first device, then press and hold the **Shift** key while clicking the last device.
- 5 Click **Save** to create the new device group, or **Cancel** to abandon it.

To edit a device group:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 1 Click **Devices** in the sub-menu, and then click the **Device Groups** tab.
- 2 Locate the device group that you want to edit and click the Edit icon () on that row.
- 3 Change the Name, add, or remove devices from the selection. Ctrl-Click devices that are not selected to select them, or Ctrl-Click selected devices to remove them from the selection.
- 4 Click **Save** to update the device group, or **Cancel** to abandon your changes.

To delete a device group:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Devices** in the sub-menu, and then click the **Device Groups** tab.
- 3 Locate the device group that you want to delete and click the delete icon () on that row. Deleting a device group does not affect the set of devices.
- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the device group.

Event Archives

Event Archives enable you to save the events for any day in the past, *not including the current day*. Archive Storage Settings must be configured before Event Archives can be created. Archive Storage Settings specify the location to which event archives will be written.



Ensure that both Configuration Backups (for configuration settings) and Event Archives (for data) run on a regular basis and are stored in a remote location. In the event of catastrophic failure, you will need to restore the most recent Configuration Backup and Event Archive. For information on Configuration Backups, see [“Configuration Backup and Restore” on page 330](#).

- For Logger appliances, the location needs to be an NFS mount, CIFS mount, or SAN, which is configured using the Logger user interface.
- For software Loggers, the location is a directory (either local or a mount point that you have already established on the machine on which the Logger software is installed).

Events in each storage group are archived separately. That is, one archive file is created for each storage group, for each day. In addition, you can bulk archive events—that is, specify a range of dates to archive events in a single archive operation.

Archiving events from each storage group to a separate archive location enables you to keep data in specific storage groups longer than others. You need to specify these locations when you configure the Archive Storage Settings before archiving any events, as shown in the following figure.



The above figure is from a Logger appliance. The Mount Location field is not available on a software Logger.

- For Logger appliances, the path you specify in the Archive Path field is appended to the path specified in the Mount Location.
- On a software Logger, you need to enter a complete path where the archive file will be written in the Archive Path field. This path could be a local directory or a mount point that is already established on the machine on which the Logger software is installed. The Mount Location field is not available on a software Logger.

Logger uses the receipt time of an event to determine its archival day. For example, an event with timestamp of 11:55:00 p.m. on December 7 is received at 12:01:00 a.m. on December 8 on the Logger. This event is archived in the archive file created for December 8th and not December 7th. When an archive operation occurs, one archive file per storage group is created at the location specified in Archive Storage Settings. Each archive file contains events from 12:00:00 a.m. to 11:59:59 p.m. for a single storage group of any given day. When you specify a range of dates, one archive file per storage group, for each specified day is created.

You can archive events in two ways: manually and scheduled. When archiving events manually, you specify the start and end dates of the event archive, and the storage groups that should be archived. This operation occurs one-time, for the specified date range. When scheduling event archives, you specify the time at which the archive operation should occur every day and select the storage groups that should be included.



You cannot set event archives to start at 1 a.m. for scheduled archives. This restriction is by design to account for the Daylight Savings Time (DST) changes.

When Logger starts archiving, it proceeds sequentially through the various storage groups, as listed on the Daily Task Settings page (for scheduled archives) or the Add Event Archives page (for manual archives).

Once the events have been archived, they are not deleted from the local storage until the events (and their related indexing information) age out due to the configured retention policy. These events continue to be included in search operations until they age out.

Once events that have been archived are deleted from Logger's local storage, they are not included in search operations. To include such events in search operations, you must load the archive in which those events exist back to the Logger. When an Event Archive is loaded, its events are included in searches, but the archive itself remains on the remote storage.

When events are archived, index information for those events is not archived. Therefore, when event archives are loaded, indices are not available. As a result, a search query that runs on archived events (that have been loaded on Logger) is slower than when the data was not archived because the index data for the archived data is not available.

The source type information (if associated with an event) is preserved when the event is archived. For information on creating and using source types, see ["Source Types" on page 274](#).

Guidelines for Archiving Events

- Be sure to run Configuration Backups as well as Event Archives regularly, and to store them in a remote location. In the event of catastrophic failure, you will need to restore the most recent Configuration Backup and Event Archive. For information on Configuration Backups, see ["Configuration Backup and Restore" on page 330](#).
- If you need to archive a large number of events (in the order of tens of GB), HP recommends that you archive during the off-peak hours to prevent impacting the performance of your Logger.
- Multiple archiving operations such as loading, unloading, archiving, and deletion of archives can occur simultaneously. Therefore, you can initiate the loading of an existing archive, while an archive operation is in progress.
- Only one manual archive job can run at a time. However, a scheduled archiving operation can run in parallel with a manual job.
- You cannot re-archive the events that have been archived already. If you try to do so, the Logger reports an error.
- Do not move the archived files from their archive location. The archives that have been moved from the originally archived location cannot be loaded on to the Logger. If you need to delete the archives, use the Logger user interface to do so.
- If an archive job fails, you need to initiate it manually. To do so, delete the failed archive and archive it manually. To get notified of a failed archive, configure an alert for this audit event: Event Archive Failed. For more information about this event, see [Appendix C, Logger Audit Events, on page 559](#). For more information about configuring alerts, see ["Alerts" on page 296](#).
- If a Logger appliance goes down while an archive operation is in progress, you need to re-initiate the archive operation for only the storage groups that were not archived when the operation failed. The status of such storage groups is marked "Failed" in the Status column on the Event Archives page. For example, you archive the event data of 12/1/10, which consists of events from four storage groups "Default", "Internal", "Short-Term", and "Long-Term". The appliance goes down after the events from the "Default" and "Internal" groups have been successfully archived, and the events from "Short-Term" are being archived. The status of the "Short-Term" storage group on the Event Archives page will display "Failed", while the status of the "Default" and "Internal" groups will display "Archived". (The status of the "Long-Term" storage group

will not be displayed.) In this case, you need to manually re-initiate the archive for the “Short-Term” and “Long-Term” storage groups.



In the above example, the status of the “Long-Term” storage group is not displayed on the Event Archives page after the failure occurs because archival of this group was never initiated during that archive operation.

If an archive operation fails, make sure you determine the storage groups that could not be archived and re-initiate the archival for all of those groups manually.

- You can cancel an in-progress archive operation that was manually initiated at any time using the Cancel link that displays on top of the Event Archives page.
- If you are upgrading from 5.0 Patch 2 or earlier, you need to note the following changes:
 - ◆ The existing event archives **cannot** be converted to make use of the storage-group level granularity that 5.1 or later releases offer. However, any data archived after you upgrade to 5.1 or later, will be archived using the storage-group level granularity. Therefore, **after upgrading to 5.1 or later**, specify the archive locations for each storage group on the Archive Storage Settings page. By default, the location you had configured-* before the upgrade is used for all storage groups.
 - ◆ Starting with 5.1, the archive locations specified on the Archive Storage Settings page can be changed anytime unlike the one-time configuration that was possible prior to this release.
 - ◆ If you change the archive location, the archives that were created on the previously configured location cannot be moved to the new location.
 - ◆ The Logger user interface for Event Archives was updated to display relevant information for the archiving changes introduced in 5.1. For example, prior to 5.1, only the name and date of the archive was displayed. However, starting with 5.1, the name, date, and storage group name are displayed—the name of the storage group to which the archives pertains is displayed as a separate column on the Event Archives page, as shown in the following figure.

Archiving Events

To save events for a particular day, you need to add an Event Archive. The table in the Event Archives tab shows the current archives and their status.

An archive storage location must be established on the Logger before you can archive its events. This is a one-time configuration. To establish an archive storage location, see [“Archive Storage Settings” on page 255](#).

To add an Event Archive:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.

Event Archives Daily Task Settings Archive Storage Settings								
Add Remove Load Unload Refresh								
<input type="checkbox"/>	Name Filter: All	Day Filter: All	Month Filter: All	Year Filter: All	Storage Group Filter: All	Status Filter: All	Mount Filter: All	Mount Path Filter: All
<input type="checkbox"/>	archive [2012-11-28] [Default Storage Group]	28	11	2012	Default Storage Group	Archived	test1	
<input type="checkbox"/>	archive [2012-11-28] [StorageGroup2]	28	11	2012	StorageGroup2	Archived	test1	
<input type="checkbox"/>	archive [2012-11-28] [StorageGroup3]	28	11	2012	StorageGroup3	Archived	test1	
<input type="checkbox"/>	archive [2012-11-28] [StorageGroup4]	28	11	2012	StorageGroup4	Archived	test1	
<input type="checkbox"/>	archive [2012-11-28] [StorageGroup5]	28	11	2012	StorageGroup5	Archived	test1	
<input type="checkbox"/>	archive [2012-11-27] [Default Storage Group]	27	11	2012	Default Storage Group	Failed	test1	

- Click **Add** in the Event Archives tab, in the right panel.
- Enter a meaningful name in the Name field for the new Event Archive and specify the Start and End dates in the format m/dd/yy, where m is month number, dd is the day of the month (with a leading zero if necessary), and yy is the two-digit year number.

When the Start and End dates are different, one archive file per storage group, for each specified day is created. For example, if you specify the following Start and End dates:

Start Date: 8/12/12

End Date: 8/13/12

And, you select two storage groups—Internal Event Storage Group and Default Storage Group. Then, four archive files will be created as a result of this archive operation—two files per storage group for the specified two days.

The Event Archives table (under the Event Archives tab) lists the archives by an alias in this format: *<archive_name> [<yyyy-m-dd>] [<storage_group_name>]*.

- Select the names of storage groups that need to be included in the archive.
- Click **Save** to start archiving events, or **Cancel** to quit.



You can cancel an in-progress archive operation at any time using the Cancel link that displays on top of the Event Archives page.

Note

To delete an Event Archive:

- Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- Click **Event Archives** in the left panel.
- Click the checkboxes in the left-most column to select the event archives that you want to delete.
- Click **Remove** from the top of the screen to delete the selected archives.
- Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Event Archive.

Scheduled Event Archive

You can schedule a daily event archive and specify what hour of the day it should run. Scheduled event archives that have finished running appear on the archive list on the Event Archives tab. Only one scheduled event archive can run at a time; however, it can run in parallel with a manually scheduled archive.

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 363](#) before you schedule an event archive.

To schedule a daily event archive:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.
- 3 Click the **Daily Task Settings** tab in the right panel.
- 4 Select a time from the “Time For Daily Archive to Start” list. Scheduled archives must start on the hour, and midnight and 1:00 AM are not on the list to allow your Logger to receive all of the previous day’s events.
- 5 Select the storage groups whose events should be included in the scheduled archive.
- 6 Click **Save** to schedule daily event archive, or click on another tab or sub-menu to cancel.

Archive Storage Settings

On the Logger appliance, Event Archives are saved to a specific NFS or CIFS mount point, or SAN. For the software Logger, event archives are saved to the specified directory, which can be a path to a local directory or to a mount point on the machine on which the software Logger is installed. To establish a mount point, see your system’s operating system documentation.

To perform Archive Storage Setting setup:

- 1 If you are using the Logger appliance, create the NFS or CIFS mount point. (See [“Storage” on page 372](#) and [“Remote File Systems” on page 372](#).)

If you are using the software version of Logger and intend to use an NFS or CIFS mount point, ensure that the external storage point is mounted on the machine on which Logger is installed. See your system’s operating system documentation for more information.
- 2 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 3 Click **Event Archives** in the left panel.
- 4 Click the **Archive Storage Settings** tab in the right panel.
- 5 Specify a mount location and an archive path for each storage group. You can specify a different path for each storage group, thus enabling the Logger to archive events to a different location for each storage group.

For Logger appliances, choose the name of an NFS mount, CIFS mount, or SAN mount point for the Mount Location field. This drop-down list contains the names you specified when creating the NFS, CIFS, or SAN mount points (**System Admin > NFS/CIFS/SAN**).

For software Loggers, the Mount Location field does not exist. You need to enter a complete path where the archive file will be written in the Archive Path field. This path could be a local directory or a mount point that is already established on the machine on which the Logger software is installed.



You must configure settings for all storage groups on the Archive Storage Settings tab even if you do not intend to archive all of them.

The above figure shows a screen from the Logger appliance. On a software Logger, the Mount Location field is not available.

- 6 Click **Save**.

Loading and Unloading Archives

Archived events must be loaded back on Logger before they can be included in a search operation. When an Event Archive is loaded, its events are included in searches, but the archive itself remains on the remote storage. When an Event Archive is unloaded, it is available for loading, but its events are not included in searches. You can unload a loaded archive if you no longer need to include it in your search operations.

To load or unload an Event Archive:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Archives** in the left panel.
- 3 Click the checkboxes in the left-most column to select the event archives that you want to load or unload.
- 4 Click **Load** or **Unload** from the top of the screen to load or unload the selected archives.

Storage

Different storage groups support the implementation of multiple retention policies. Each group can have a different policy, and storage rules determine which storage group is used for events from specific device groups. See [“Retention Policy” on page 30](#). The Storage section has three tabs: Storage Groups, Storage Rules, and Storage Volume.



Storage Groups					
Storage Rules					
Storage Volume					
Name	Maximum Age (Days)	Maximum Size (GB)	Creator	Last Editor	
Default Storage Group	180	385	admin	admin	
Internal Event Storage Group	365	5	System	System	

Figure 6-2 Storage Groups page

Storage Groups

Storage Groups support multiple retention policies by defining a maximum size (Maximum Size) and number of days (Maximum Age) to retain events. Once events are older than the specified Maximum Age or there are more events than the storage group will hold (as specified by Maximum Size), the oldest events are deleted at the next retention cycle. The retention process triggers periodically on Logger, therefore, events might not be deleted immediately when events gets older than Maximum Age or the storage group size exceeds the Maximum Size limits.


Logger can have a maximum of 6 storage groups—two that pre-exist on your Logger (Internal Storage Group and Default Storage Group) and four that you can create. You can add the additional storage groups (up to the maximum of six) at any time.

HP recommends that you create the four additional storage groups in addition to the two that preexist, so that you have five storage groups available for event storage and one for Logger's internal events.

To add additional storage groups, follow the instructions in [“Adding Storage Groups” on page 343](#).

Once a storage group is created, it cannot be deleted however its size can be increased or decreased any time. If you are decreasing the size of the storage group and the new size is lesser than the currently used space on the storage group, you will need to delete data to achieve the new size. Logger UI guides you in this situation to delete sufficient data.

To edit (including resizing) a storage group:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Storage** in the left panel. The Storage Groups tab displays the available storage groups.
- 3 Identify the storage group you want to modify and click the Edit icon () for it. The Storage Groups tab displays the Edit Storage Group pane.

Storage Groups Storage Rules Storage Volume

Edit Storage Group

Important: Please be aware that if you reduce the maximum age from the current value you will lose all events that are older than the new maximum age.

Reducing the maximum storage group age may take a few minutes to complete.

Maximum Age (Days) 365

Maximum Size (GB) 5

Current Size (GB) 1

Save Cancel

- 4 Change the desired parameters such as the name of the storage group, or increase or decrease Maximum Age or Maximum size.



The names of the Internal Storage Group and Default Storage Group cannot be modified.

If you are reducing the size of the storage group and the new size is smaller than the value indicated by the Current Size field on the Edit Storage Group page, Logger displays a message indicating that reducing storage group size in this situation will require you to delete existing data.

If you choose to delete data to reduce the storage group size, follow these steps:

- a Set the Maximum Age value to the number indicated in the above message. Doing so, triggers deletion of events.
- b Refresh the Edit Storage Group screen. When the "Current Size" value is less than or equal to the storage group size you want to set, go to the next step. Otherwise, keep refreshing the screen periodically.



The "Current Size" value changes as data are deleted, which can take some time. Therefore, you need to wait before proceeding to the next step.

- c Set the Maximum Size value to suit your needs.
- d If you wish, restore the Maximum Age setting (that you changed in Step b) to the original value.

If you choose **not** to delete data, go to the next step to exit the procedure.



If there is sufficient space to reduce the storage group size, you can change it without modifying the Maximum Age value (to modify the retention policy to delete data).

- 5 Click **Save** to store the changes, or **Cancel** to quit.

Storage Rules

Storage rules create a mapping between device groups and storage groups. Doing so enables you to store events from specific sources to a specific storage group. Additionally, you can configure these storage groups with different retention policies, and thus retain event data based on the source of incoming events. For example, all events from firewall devices can be subject to a short retention period. To accomplish this, manually assign the firewall devices to a device group and then create a storage rule that maps the device group to a storage group with the desired short retention period.



Events that are not subject to any storage rule are sent to the Default Storage Group.

Before you add a storage rule, make sure that the storage group to which you want to store the events and the device group that contains the devices whose events you want to store exist. For information on how to create device groups, see [“Device Groups” on page 249](#).

Logger allows you to create up to 40 storage rules. If you create additional rules, an error might be generated.

To add a storage rule:


- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Storage** in the left panel, and then click the **Storage Rules** tab.
- 3 Click **Add**. The following page is displayed.

- 4 Enter the following parameters:


Parameter	Description
Storage Group	Select a storage group from the drop-down list. The storage groups must already be set up before any storage rules are added. You can only add storage groups at the time of Logger initialization.
Device Groups	Select one or more device groups to associate with the specified storage group. You may associate several device groups with a single storage group.
Priority	An integer that indicates the new rule's priority. The number must be unique for each storage rule. The smaller the number, the higher the rule's priority.

- 5 Click **Save** to add the new storage rule, or **Cancel** to quit.

To edit or reorder a storage rule:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Storage** in the left panel.
- 1 Click the **Storage Rules** tab.
- 2 Find the storage rule that you want to edit and click the Edit icon () on that row.
- 3 Change the information in the form--for example, change the priority value to reposition the storage rule in the table--and click **Save**.

To delete a storage rule:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Storage** in the left panel.
- 3 Click the **Storage Rules** tab.
- 4 Find the storage rule that you want to delete and click the Remove icon ().
- 5 Click **OK** to confirm the delete.

Storage Volume

The Storage Volume tab displays the mount location and current storage volume settings. To increase the Storage Volume size, go to the System Maintenance page. You must have admin-level privileges to perform this operation. For more information, see [“Storage Volume Size Increase” on page 342](#).



Storage volume can be extended after initialization, but its size cannot be reduced.

Note

Event Input

The event input section allows you to set up how the data that comes into Logger is described. You can configure receivers to listen for and capture event data, set up source types that identify the type of log file that captured events come from, and create parsers that extract field-value pairs when searching the data.

Receivers

Logger can receive text events, either sent through the network or read from a file. From the Receivers tab, you can set up and configure the receivers that will capture event data for your Logger, and populate each event with information about its origin. Some receivers capture streaming events transmitted over the network by devices, applications, services, and so on. Other types of receivers monitor individual files for events or monitor files selected from a directory tree, based on a pattern you specify. Since receivers can only receive events of a single source type, you should set up separate receivers for each type of log file. To start receiving events, direct your event sources to the default receivers. For more information about the default receivers, see [“Receivers” on page 32](#).

Receiver types include UDP, TCP, SmartMessage, and three types of file based receivers, File Transfer, File Receiver, and Folder Follower Receiver.

You can configure the following types of receivers for Logger:

- **UDP Receiver:** UDP receivers listen for User Datagram Protocol messages on the port you specify. Logger comes pre-configured with a UDP Receiver on port 514 or 8514, enabled by default. For software Loggers, this port may vary based on the port numbers available at installation time.
- **CEF UDP Receiver:** UDP receivers that receive events in Common Event Format.
- **TCP Receiver:** TCP receivers listen for Transmission Control Protocol messages on the port you specify. Logger comes pre-configured with a TCP receiver on port 515 or 8515, enabled by default. For software Loggers, this port may vary based on the port numbers available at installation time.
- **CEF TCP Receiver:** TCP receivers that receive events in Common Event Format.
- **File Receiver:** Depending on the type of Logger, file receivers read log files from a local file system, Network File System (NFS), Common Internet File System (CIFS), or Storage Area Network (SAN). File receivers read single or multi-line log files. They provide a snapshot of a log file at a single point in time.
- **Folder Follower Receiver:** Folder follower receivers actively read the log files in a specified directory as they are updated. If the source directory contains different types of log files, you can create a receiver for each type of file that you want to monitor. Logger comes pre-configured with folder follower receivers for Logger's Apache Access Error Log, the system Messages Log, and Audit Log (when auditing is enabled). You must enable these receivers in order to use them.
- **File Transfer:** File Transfer receivers read remote log files using SCP, SFTP or FTP protocol. These receivers can read single- or multi-line log files. You can schedule the receiver to read a file or batch of files periodically.



- The SCP, SFTP, and FTP file transfer receivers depend on the FTP (File Transfer Protocol) SCP (Secure Copy Protocol) and SFTP (SSH file transfer protocol) clients installed on your system. Before creating the receiver, ensure that the appropriate client is installed on the system.
- The SCP and SFTP protocols on Logger appliances are not FIPS compliant.

- **SmartMessage Receiver:** SmartMessage receivers listen for encrypted messages from ArcSight SmartConnectors. Logger comes pre-configured with a SmartMessage receiver with the name "SmartMessage Receiver". To use this receiver to receive events from a SmartConnector, set the **Receiver Name** to be "SmartMessage Receiver" when configuring the SmartConnector's destination.

File Based Receivers

File based receivers include File Receiver, File Transfer Receiver and Folder Follower Receiver. You can set them up as multiline receivers, and configure them to use source types with associated parsers to extract data from captured events.



When a receiver cannot read the file it logs from, such as when the file or folder is deleted or renamed, logger records a message in `current/arcSight/logger/logs/Logger_receiver.log`.

Multi-line Receivers

TCP and UDP receivers interpret line break characters, such as `\r` or `\n`, as the end of the event. If the input event contains embedded `\r` or `\n` characters, the event will be treated as more than one event. If your events span more than one line, you may want to use a multi-line receiver. Multi-line receivers include the File Transfer, File Receiver and Folder Follower Receivers.

A multi-line receiver can read events that span more than one line, such as a server log. You could set up the receiver to handle stack traces reported in the log by reading the entire stack trace as a single event instead of reading each line separately.

When creating a multi-line receiver, you must specify a regular expression that the receiver should use to detect the start of a new event in the log file. Each new event starts where the characters in the log file match the regular expression.

For example, in the following log file, each event starts with a timestamp embedded within square brackets (`[yy-MM-dd HH:mm:ss.SSS]`); therefore, you can use this regular expression to identify each event:

```
^\[\d+-\d+-\d+ \d+:\d+:\d+,\d+\].*
```

```
[2010-12-06 13:11:26,824][INFO ][I18N]Locale has not been chosen by the user.
[2010-12-06 13:11:26,828][ERROR][DirectConnection$ReadChannel]
java.io.IOException: end of communication channel
    at com.arcsight.logger.distributed.DirectConnection.a(DirectConnection.java:39)
    at com.arcsight.logger.distributed.DirectConnection.access$200(DirectConnection.java:19)
    at com.arcsight.logger.distributed.DirectConnection$ReadChannel.run(DirectConnection.java:85)
    at java.util.concurrent.ThreadPoolExecutor$Worker.runTask(ThreadPoolExecutor.java:886)
    at java.util.concurrent.ThreadPoolExecutor$Worker.run(ThreadPoolExecutor.java:908)
    at java.lang.Thread.run(Thread.java:619)
```

- For multi-line file receivers and file transfer receivers, the regular expression that identifies the beginning of a new event must be specified in the receiver's *Multiline Event Starts With* field.
- For multi-line folder follower receivers, the regular expression that identifies the beginning of a new event must be specified in the *Multiline Event Starts With* field of the **source type** associated with that receiver, rather than in the receiver itself.

For information on creating and using receivers, see [“Working with Receivers” on page 263](#). For information on creating and using source types, see [“Source Types” on page 274](#).

Folder Follower Receivers

When you want to monitor active files as they are updated, use a folder follower receiver. After you set up a folder follower receiver and enable it, it will monitor the specified files in that directory and continuously upload new events to Logger. Folder follower receivers recognize file rotation.

Overview of the steps to monitor a directory:

- 1 Determine the types of logs you need to monitor.
- 2 Determine whether the out-of-box source types or source type/parser pairs will satisfy your needs. For more information, see [“Source Types” on page 274](#), and [“Parsers” on page 278](#).
 - ◆ If so, proceed to [Step 3](#).

- ◆ If not, create the parsers and sources types that you need.
 - i Select an appropriate parser or set of parser for the log files in the directory you want to follow. If the out-of-box parsers do not provide what you need, create appropriate parsers.
 - ii Assign a source type for each parser. If the out-of-box source types do not provide what you need, create appropriate source types.
- 3 Create the folder follower receivers required to monitor the logs in the directory, selecting the source type determined in [Step 1](#). For more information, see [“Working with Receivers” on page 263](#).
- 4 Enable the receivers.
- 5 Optionally, to forward log file events, set up and configure one or more forwarders. For more information, see [“Forwarders” on page 285](#).

Using Source Types with File Follower Receivers

Logger uses the parser associated with the source type you select for a receiver to extract fields and their respective values from the received events. These fields are parsed at search time. For more information on using source types and parsers, see [“Source Types” on page 274](#), and [“Parsers” on page 278](#).

When creating a file follower receiver, you must select a source type appropriate to monitor a specific type of log file. After you select the source type for the file follower receiver, ensure that the parser associated with it works with your source files.

Events from different versions of the same source type can be in different formats. Similarly, events from different source types of the same vendor might be formatted differently. Therefore, if the source type you choose from the Logger UI does not exactly match the specifications of your source type, the associated parser will not parse events correctly, and the search results will not display any parsed fields.

To confirm whether the source type has a valid parser for your source type, after you have set up the receiver, check whether the incoming events are parsed. To determine this, run a search and review the “parser” field in the search results. The parser used in the search will be displayed in the parser column of the search results. If the event was parsed, this field contains the name of the parser. If the event was not parsed successfully, this field contains “Not parsed”. If no parser is defined for the source type or if there is no source type, the field is blank.

Working with Receivers

Several receivers come set up on your Logger. You can add other receivers as needed. The maximum number of receivers that can be created on Logger is limited by system resources—memory, CPU, disk input/output and possibly network bandwidth.

The receiver ports available on your Logger may vary from the image below.

Receivers

Source Types

Parsers

Add

Once you enable the Apache URL Access Error Log receiver, Logger will start storing entries from the <install_dir>/userdata/logs/apache/http_error_log file.

Logger can also store entries from the messages and audit.log files in the /var/log/* folders. Before enabling the receivers for these files, consult the Logger Administrators guide for details.

Name	Type	IP Address	Port			
Apache URL Access Error Log	Folder Follower Receiver					
Audit Log	Folder Follower Receiver					
Var Log Messages	Folder Follower Receiver					
SmartMessage Receiver	SmartMessage Receiver					
TCP Receiver	TCP Receiver	All	515			
UDP Receiver	UDP Receiver	All	514			

Figure 6-3 Receivers tab



Note

Before creating a receiver of type File Receiver:

- For the Logger appliance, set up a Network File System mount. See [“Storage” on page 256](#).
- For the software version of Logger, the file system from which the log files will be read needs to be mounted on the system on which you have installed Logger.

Before creating a receiver of type File Transfer, ensure that the appropriate SCP, SFTP, and FTP client is installed on your system.

To create a receiver:

- Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- Click **Event Input** (left panel) > **Receivers** tab (right panel).
The Receivers tab, shown in [Figure 6-3 on page 264](#), displays the current receivers and their status. You can sort the fields by clicking the column headers.
- Click **Add**.
- Enter a name for the new receiver. SmartMessage receiver names are used when configuring the associated ArcSight SmartConnectors.
- Choose the receiver type. Select UDP Receiver, TCP Receiver, CEF UDP Receiver, CEF TCP Receiver, File Receiver, Folder Follower Receiver, File Transfer, or SmartMessage Receiver.



Note

The receiver type cannot be changed after the receiver is created.

- Click **Next** to edit receiver parameters.
The fields displayed in the Edit Receiver dialog box vary according to the type of Logger and the type of receiver.
- Fill in the appropriate fields. Refer to the following tables for field descriptions.

- ◆ [Table 6-1, "Parameters used in UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers," on page 266](#)
- ◆ [Table 6-2, "Parameters used in File Receivers," on page 267](#)
- ◆ [Table 6-3, "Parameters used in Folder Follower Receivers," on page 269](#)
- ◆ [Table 6-4, "Parameters used in File Transfer Receivers," on page 270](#)
- ◆ [Table 6-5, "Parameters used in SmartMessage Receivers," on page 273](#)

8 Click **Save**.

9 New receivers are initially disabled. You must enable them in order to use them.

To enable or disable a receiver:



Note

Before enabling the following preconfigured folder follower receivers for the software version of Logger, ensure that the files are readable by the non-root user that you installed with or specified during installation.



- `/var/log/messages`
- `/var/log/audit/audit.log`

1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.

2 Click **Event Input** (left panel) > **Receivers** tab (right panel).

The Receivers tab, shown in [Figure 6-3 on page 264](#), displays the current receivers and their status. You can sort the fields by clicking the column headers.

3 Locate the receiver that you want to enable or disable.

- ◆ If the receiver is currently disabled, click the disabled icon () to enable it.
- ◆ If the receiver is currently enabled, click the enabled () icon to disable it.



Tip

Wait a few minutes after enabling a receiver before disabling it. Likewise, wait before enabling a receiver that has just been disabled. Background tasks initiated by enabling or disabling a receiver can produce unexpected results if they are interrupted.

To edit a receiver:

1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.

2 Click **Event Input** (left panel) > **Receivers** tab (right panel).

The Receivers tab, shown in [Figure 6-3 on page 264](#), displays the current receivers and their status. You can sort the fields by clicking the column headers.

3 Locate the receiver that you want to update and click the Edit icon () on that row.

The fields displayed in the Edit Receiver dialog box vary according to the type of Logger and the type of Receiver.

4 Edit the appropriate fields. Refer to the following tables for field descriptions.

- ◆ [Table 6-1, "Parameters used in UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers," on page 266](#)
- ◆ [Table 6-2, "Parameters used in File Receivers," on page 267](#)
- ◆ [Table 6-3, "Parameters used in Folder Follower Receivers," on page 269](#)

- ◆ [Table 6-4, “Parameters used in File Transfer Receivers,” on page 270](#)
- ◆ [Table 6-5, “Parameters used in SmartMessage Receivers,” on page 273](#)

5 Click **Save**.

To delete a receiver:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Receivers** tab (right panel).

The Receivers tab, shown in [Figure 6-3 on page 264](#), displays the current receivers and their status. You can sort the fields by clicking the column headers.

- 3 Locate the receiver that you want to delete and click the Remove icon (✖) on that row.
- 4 Click **OK** to confirm the delete.

Receiver Parameters

The following tables describe the parameters used when creating and editing receivers:

- [Table 6-1, “Parameters used in UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers,” on page 266](#)
- [Table 6-2, “Parameters used in File Receivers,” on page 267](#)
- [Table 6-3, “Parameters used in Folder Follower Receivers,” on page 269](#)
- [Table 6-4, “Parameters used in File Transfer Receivers,” on page 270](#)
- [Table 6-5, “Parameters used in SmartMessage Receivers,” on page 273](#)

Fill in the following fields when creating or editing UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers.

Table 6-1 Parameters used in UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers

Parameter	Description
Name	The name of the receiver, used for reporting and status monitoring.
IP/Host	Select one of the Logger's network connections for the receiver to listen to, or select All to listen on both network connections. Note: If localhost (127.0.0.1) appears in the list, it means that the Logger hostname has not been configured. To configure hostname, see “Network” on page 358 .
Port	For the Logger appliance: The default UDP Receiver is pre-configured on port 514. For SmartMessage receivers, configure the SmartConnector for port 443. For the software version of Logger: If you installed software Logger as a root user, you can use any available port. The default UDP Receiver is pre-configured on port 514. If that port is not available, then the next higher available port is chosen. If you installed software Logger as a non-root user, you can only use a port numbers greater than 1024. The default UDP Receiver is pre-configured on port 8514. If that port is not available, then the next higher available port is chosen.

Table 6-1 Parameters used in UDP Receivers, TCP Receivers, CEF UDP Receivers, and CEF TCP Receivers (Continued)

Parameter	Description
Encoding	Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.
Source Type	<p>Select from the pulldown list of log file types, including:</p> <ul style="list-style-type: none"> • Apache HTTP Server Access • Apache HTTP Server Error • Juniper Steel-Belted Radius • Microsoft DHCP Log • IBM DB2 Audit <p>Additionally, you can define your own source type, based on the needs of your company. See “Source Types” on page 274.</p> <p>A receiver can only receive events of a single source type. Set up separate receivers for each type of log file.</p> <p>Note: CEF TCP and CEF UDP receivers are set to the CEF source type, and cannot be changed. Currently, there is no parser associated with the CEF source type.</p> <p>Associating Source types with TCP and UDP receivers was introduced in Logger 5.3 SP1. When upgrading, TCP and UDP receivers from earlier releases are set to the “Other” source type.</p>

Fill in the following fields when creating or editing File Receivers.

Table 6-2 Parameters used in File Receivers

Parameter	Description
Name	The name of the receiver, used for reporting and status monitoring.
RFS Names	<p>Select from the pulldown list of NFS or CIFS mount names. The list also includes attached SANs on Logger models that support SAN.</p> <p>To mount NFS volumes, see “Storage” on page 372. To mount CIFS shares, see “Remote File Systems” on page 372. For more information about SAN, see “SAN” on page 375.</p>
Folder	<p>Choose “Local” and then specify the directory on your Logger where the remote file system is mounted in the “Folder” field.</p> <p>To mount a remote file system on the system on which you have installed Logger, see its operating system’s documentation.</p>
Source Type	<p>Select from the pulldown list of log file types, including:</p> <ul style="list-style-type: none"> • Apache HTTP Server Access • Apache HTTP Server Error • Juniper Steel-Belted Radius • Microsoft DHCP Log • IBM DB2 Audit <p>Additionally, you can define your own source type, based on the needs of your company. See “Source Types” on page 274.</p> <p>A receiver can only receive events of a single source type. Set up separate receivers for each type of log file.</p>

Table 6-2 Parameters used in File Receivers (Continued)

Parameter	Description
Wildcard (regex)	<p>A regular expression (regex) describing the log files to read.</p> <p>This is a regular expression, not a typical file wildcard like <code>"*.*</code>".</p> <p>The default is <code>.*</code>, meaning all files.</p> <p>Examples:</p> <p>To include all files ending with <code>.process</code>, you could use:</p> <pre>.*\.process</pre> <p>To monitor only <code>*.properties</code> files, you could use:</p> <pre>.*\.properties</pre> <p>To include only <code>.log</code> files with eight digit filenames, you could use:</p> <pre>\d{8}\.log</pre> <p>Note: Uploading any type of data other than text, including binary files such as <code>.zip</code> or <code>.bin</code>, may prevent Logger from functioning correctly. Use caution when pulling everything from a directory by specifying <code>.*</code> in the Regex field, as you could inadvertently include binary files.</p>
Mode	<p>Select one of the following:</p> <p>Delete - delete the log file once it has been processed</p> <p>Rename - rename the log file once it has been processed. The file is named by appending the Rename Extension.</p> <p>Persist - Logger remembers which files have been processed and only processes them once.</p>
Rename extension	The suffix to append to log files that have been processed.
Character encoding	Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.
Delay after seen	<p>Number of seconds to wait after a source file is first seen until it is processed. This allows the entire file to be copied to Logger or (in the case of File Receiver) copied to the remote file system, before processing begins.</p> <p>The default is 10 seconds.</p> <p>Note: For File Transfer Receivers, this parameter should be set to a larger value if large files are expected. The default, 10 seconds, does not allow enough time for a large file, such as 1 GB.</p>
Event Time Locale	Select a locale from the pulldown list, such as English (United States), Chinese (Hong Kong), Chinese (Taiwan), and so on.
Date/time zone	<p>Required if the timestamp in the log file does not specify a time zone.</p> <p>For File Transfer and File Receivers, this parameter is ignored if either Date/time format or Date/time location regex are blank.</p> <p>You can see the time zone configured on the Logger (System Admin > Settings > Platform > Time/NTP).</p> <p>Software Loggers use the system time.</p>
Event Time Location	<p>A regular expression describing which characters represent the timestamp in the log file. For example:</p> <pre>.*\[(.*)\].*</pre> <p>This regular expression specifies that the timestamp is found inside the first set of square brackets on each line. The first capturing group (the part of the regex in parentheses) is that part that is then parsed using the Date/time format.</p> <p>The default is <code>"</code> (no timestamp in log file).</p>

Table 6-2 Parameters used in File Receivers (Continued)

Parameter	Description
Event Time Format	<p>Required if the log file contains timestamps in the same format for each event. If not specified (or if the Date/time location regex is blank), each event in the file will be stamped with the date that the file itself was first seen by Logger (not its file system timestamp).</p> <p>See Step Table 6-7 for a list of format specifiers.</p> <p>The default is "" (no timestamp in log file).</p>
Multiline Event Starts With	<p>A regular expression that specifies the start of a new event in a log file. Specify this expression to enable the receiver to read multi-line log files. Each new event starts at the point where the regular expression is matched to the characters in the log file. For example,</p> <pre>^\[\d+-\d+-\d+ \d+:\d+,\d+\].*</pre> <p>This regular expression matches timestamps such as:</p> <pre>[2010-12-06 13:09:46,818]</pre> <p>When this field is left blank (""), each line in the log file is treated as a single event.</p> <p>The default is "" (each line in the log file is a single event).</p>

Fill in the following fields when creating or editing Folder Follower Receivers.

Table 6-3 Parameters used in Folder Follower Receivers

Parameter	Description
Name	The name of the receiver, used for reporting and status monitoring.
Local Folder	Specify the local folder to process. On the Logger appliance, this field is only available if you select "Local" for the Mount Name.
Source Type	<p>Select from the pulldown list of log file types, including:</p> <ul style="list-style-type: none"> • Apache HTTP Server Access • Apache HTTP Server Error • Juniper Steel-Belted Radius • Microsoft DHCP Log • IBM DB2 Audit <p>Additionally, you can define your own source type, based on the needs of your company. See "Source Types" on page 274.</p> <p>A receiver can only receive events of a single source type. Set up separate receivers for each type of log file.</p>

Table 6-3 Parameters used in Folder Follower Receivers (Continued)

Parameter	Description
Wildcard (regex)	<p>A regular expression (regex) describing the log files to read.</p> <p>This is a regular expression, not a typical file wildcard like <code>"*. *"</code>.</p> <p>The default is <code>.*</code>, meaning all files.</p> <p>Examples:</p> <p>To include all files ending with <code>.process</code>, you could use:</p> <pre>.*\.process</pre> <p>To monitor only <code>*.properties</code> files, you could use:</p> <pre>.*\.properties</pre> <p>To include only <code>.log</code> files with eight digit filenames, you could use:</p> <pre>\d{8}\.log</pre> <p>Note: Uploading any type of data other than text, including binary files such as <code>.zip</code> or <code>.bin</code>, may prevent Logger from functioning correctly. Use caution when pulling everything from a directory by specifying <code>.*</code> in the Regex field, as you could inadvertently include binary files.</p>
Blacklist (regex)	<p>A regular expression (regex) describing the name of the log files to ignore. Files are not monitored if they match this expression.</p> <p>This is a regular expression, not a typical file wildcard like <code>"*. *"</code>.</p> <p>Example:</p> <p>To exclude files that end in <code>.txt</code>, you could use:</p> <pre>.*\.txt</pre> <p>To monitor all files except <code>*.txt</code>, you could use:</p> <p>Wildcard: <code>.*</code></p> <p>Blacklist: <code>.*\.txt</code></p>
Character encoding	<p>Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.</p>
Date/time zone	<p>Required if the timestamp in the log file does not specify a time zone.</p> <p>For File Transfer and File Receivers, this parameter is ignored if either Date/time format or Date/time location regex are blank.</p> <p>You can see the time zone configured on the Logger (System Admin > Settings > Platform > Time/NTP).</p> <p>Software Loggers use the system time.</p>

Fill in the following fields when creating or editing File Transfer Receivers.

Table 6-4 Parameters used in File Transfer Receivers

Parameter	Description
Name	The name of the receiver, used for reporting and status monitoring.
Protocol	Select SCP, SFTP or FTP protocol.

Table 6-4 Parameters used in File Transfer Receivers (Continued)

Parameter	Description
Port	<p>For the Logger appliance: The default UDP Receiver is pre-configured on port 514 (For SmartMessage receivers, configure the SmartConnector for port 443.)</p> <p>For the software version of Logger: If you installed software Logger as a root user, you can use any available port. The default UDP Receiver is pre-configured on port 514. If that port is not available, then the next higher available port is chosen. If you installed software Logger as a non-root user, you can only use a port numbers greater than 1024. The default UDP Receiver is pre-configured on port 8514. If that port is not available, then the next higher available port is chosen.</p>
IP/Host	<p>Select one of the Logger's network connections for the receiver to listen to, or select All to listen on both network connections.</p> <p>Note: If localhost (127.0.0.1) appears in the list, it means that the Logger hostname has not been configured. To configure hostname, see "Network" on page 358.</p>
User	A user on the host with privileges to view and read the source log files. If the protocol is FTP, you can specify the special user, "anonymous."
Password	The password of the specified User. The password must not be empty, even in the case of anonymous FTP (although in this case, the password will be ignored.)
File path	<p>The path and the name of the log file(s) to be read. You can use wild cards like ? and * (for example, "*.log" or "Log-?.txt") in the path name and the file name. Separate directories with forward slashes (/).</p> <p>Separate multiple file specifications with commas.</p> <p>Example: /tmp/SyslogData/syslog.log.gz, /security/logs/*/, /security/log?/admin/special/</p> <p>Note: Uploading any type of data other than text, including binary files such as .zip or .bin, may prevent Logger from functioning correctly. Be sure that any directories you specify do not include binary files. Use caution when pulling everything from a directory by specifying *, as you could inadvertently include binary files.</p>
Schedule	<p>If no schedule is specified, the File Transfer will occur just once.</p> <p>Choose Everyday or Days of Week from the first pulldown menu.</p> <p>If Everyday, select Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the text box.</p> <p>If Days of Week, enter the days (day 1 is Sunday) in the text box. Then choose Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the second text box.</p> <p>For example, to read log files every day at 2 a.m., select Everyday in the first pulldown menu, then choose Hour of Day from the second pulldown menu and enter 2 in the text box. To read the log files every day at 2 a.m. and 3 p.m., enter 2,15 in the text box.</p> <p>For another example, to read log files Tuesdays and Thursdays at 10 p.m., select Days of Week from the first pulldown menu and enter 3,5 for days. Then choose Hour of Day from the second pulldown menu and enter 22 in the text box.</p> <p>Make sure you are familiar with the information in "Impact of Daylight Savings Time Change on Logger Operations" on page 363 before you schedule a file transfer.</p>
Zip Format	Choose gzip, zip, or none.

Table 6-4 Parameters used in File Transfer Receivers (Continued)

Parameter	Description
Source Type	<p>Select from the pulldown list of log file types, including:</p> <ul style="list-style-type: none"> • Apache HTTP Server Access • Apache HTTP Server Error • Juniper Steel-Belted Radius • Microsoft DHCP Log • IBM DB2 Audit <p>Additionally, you can define your own source type, based on the needs of your company. See “Source Types” on page 274.</p> <p>A receiver can only receive events of a single source type. Set up separate receivers for each type of log file.</p>
Character encoding	Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.
Delay after seen	<p>Number of seconds to wait after a source file is first seen until it is processed. This allows the entire file to be copied to Logger or (in the case of File Receiver) copied to the remote file system, before processing begins.</p> <p>The default is 10 seconds.</p> <p>Note: For File Transfer Receivers, this parameter should be set to a larger value if large files are expected. The default, 10 seconds, does not allow enough time for a large file, such as 1 GB.</p>
Event Time Locale	Select a locale from the pulldown list, such as English (United States), Chinese (Hong Kong), Chinese (Taiwan), and so on.
Date/time zone	<p>Required if the timestamp in the log file does not specify a time zone.</p> <p>For File Transfer and File Receivers, this parameter is ignored if either Date/time format or Date/time location regex are blank.</p> <p>You can see the time zone configured on the Logger (System Admin > Settings > Platform > Time/NTP).</p> <p>Software Loggers use the system time.</p>
Event Time Location	<p>A regular expression describing which characters represent the timestamp in the log file. For example:</p> <pre>.*\[(.*)\].*</pre> <p>This regular expression specifies that the timestamp is found inside the first set of square brackets on each line. The first capturing group (the part of the regex in parentheses) is that part that is then parsed using the Date/time format.</p> <p>The default is "" (no timestamp in log file).</p>
Event Time Format	<p>Required if the log file contains timestamps in the same format for each event. If not specified (or if the Date/time location regex is blank), each event in the file will be stamped with the date that the file itself was first seen by Logger (not its file system timestamp).</p> <p>See Table 6-7, “Date/time format specification,” on page 274 for a list of format specifiers.</p> <p>The default is "" (no timestamp in log file).</p>

Table 6-4 Parameters used in File Transfer Receivers (Continued)

Parameter	Description
Multiline Event Starts With	<p>A regular expression that specifies the start of a new event in a log file. Specify this expression to enable the receiver to read multi-line log files. Each new event starts at the point where the regular expression is matched to the characters in the log file. For example,</p> <pre>^\[\d+-\d+-\d+ \d+:\d+,\d+\].*</pre> <p>This regular expression matches timestamps such as:</p> <pre>[2010-12-06 13:09:46,818]</pre> <p>When this field is left blank (""), each line in the log file is treated as a single event.</p> <p>The default is "" (each line in the log file is a single event).</p>

Fill in the following fields when creating or editing SmartMessage Receivers.

Table 6-5 Parameters used in SmartMessage Receivers

Parameter	Description
Name	The name of the receiver, used when configuring an associated ArcSight SmartConnectors.
Encoding	Select a character encoding, such as US-ASCII, Big5, or EUC-KR, from the pulldown list. CEF UDP, CEF TCP, and SmartMessage receivers must use US-ASCII or UTF-8 encoding.

Date and Time Specification

To specify the date and time format so that it can be parsed from a file receiver, (File Receiver, Folder Follower Receiver, or File Transfer), refer to [Table 6-7 on page 274](#).

Internally, Logger uses a common Java method called SimpleDateFormat that you may be familiar with. Sophisticated uses of SimpleDateFormat, as described in Java sources, will work with Logger. Pattern letters are usually repeated, as their number determines the exact presentation:

The examples in [Table 6-6 on page 273](#) show how date and time patterns are interpreted in the U.S. locale. The given date and time are 2001-07-04 12:08:56 local time in the "U.S. Pacific Time" time zone.

Table 6-6 Date/time examples

Source	Date and Time Pattern
2001.07.04 AD at 12:08:56 PDT	yyyy.MM.dd G 'at' HH:mm:ss z
Wed, Jul 4, '01	EEE, MMM d, ''yy
12:08 PM	h:mm a
12 o'clock PM, Pacific Daylight Time	hh 'o'clock' a, zzzz
0:08 PM, PDT	K:mm a, z
02001.July.04 AD 12:08 PM	yyyyy.MMMMM.dd GGG hh:mm aaa
Wed, 4 Jul 2001 12:08:56 -0700	EEE, d MMM yyyy HH:mm:ss Z

Table 6-6 Date/time examples (Continued)

Source	Date and Time Pattern
010704120856-0700	yyMMddHHmmssZ
2001-07-04T12:08:56.235-0700	yyyy-MM-dd'T'HH:mm:ss.SSSZ

Table 6-7 Date/time format specification

Symbol	Meaning	Presentation	Examples
G	Era designator	(Text)	AD
y	Year	(Number)	2006 or 06
M	Month in year (1-12)	(Month)	July or Jul or 07
w	Week in year (1-52)	(Number)	39
W	Week in month (1-5)	(Number)	2
D	Day in year (1-366)	(Number)	129
d	Day in month (1-31)	(Number)	10
E	Day in week	(Text)	Tuesday or Tue
F	Day in week of month		
a	Am/pm marker	(Text)	AM or PM
H	Hour in day (0-23)	(Number)	0
k	Hour in day (1-24)	(Number)	24
K	Hour in am/pm (0-11)	(Number)	0
h	Hour in am/pm (1-12)	(Number)	12
m	Minute in hour (0-59)	(Number)	30
s	Second in minute (0-59)	(Number)	55
S	Millisecond (0-999)	(Number)	978
z	Time zone	(Text)	Pacific Standard Time, or PST, or GMT-08:00
Z	Time zone	(RFC 822)	-0800 (indicating PST)

Source Types

Source types identify the kind of event that comes from a specific data source. For example, an event could come from an Apache access log, a simple syslog, or the log of an application you created. You can use parsers to parse event data from a specified source type.

Once events are associated with a source type, if the source type is associated with a parser, the events are parsed by that parser when you run a search that matches those events. The search result displays the matching parsed event fields in columns, similar to the CEF events. (Use the "User Defined Fields" field set to view these events.) For more information, see ["Parsers" on page 278](#).

The source of the event, the source type, and the parser will be displayed in the column list of the search results if any row is fetched from a search which contains a non-CEF source type.

The following columns are displayed in the search results when a source type is used:

- **Source**—The name of the log file from which the event was received. For example, `/opt/mnt/testsoft/web_server.out.log`. If no source was applied when the event was received, this field is blank. You can control whether this field is displayed from the Search Options tab. See [“Tuning Advanced Search Options” on page 321](#) for how to set this option.
- **Source Type**—The type of file from which the event was received, as defined on the Source Type page (**Configuration** (or **Configuration > Settings**) > **Event Input > Source Types**). If no source type was applied when the event was received, this field is blank. You can control whether this field is displayed from the Search Options tab. See [“Tuning Advanced Search Options” on page 321](#) for how to set this option.
- **Parser**—If the event was parsed, this field contains the name of the parser. If the event was not parsed successfully, this field contains “Not parsed”. If no parser is defined for the source type or if there is no source type, the field is blank.

Working with Source Types

Logger provides a number of source types with pre-configured parsers. Additionally, you can define new source types and assign parsers to them. This lets you choose the set of fields you want to extract for a given kind of event. Only one parser can be associated with a source type, however, multiple source types can be associated with a parser. Out-of-box source types cannot be edited or deleted, but you can copy them to make similar source types to meet your needs. You can edit or delete custom source types, as desired. The source types available on your Logger may vary from the image below.

Receivers Source Types Parsers									
Add									
Name	Description	Parser	Event Time Location	Event Time Format	Multiline Event Starts With	Locale			
test1		test1				English (United States)			
Apache_access	Apache Access Log	Apache_access	.*\[(.*)\].*	dd/MMM /yyy:HH:mm:ss Z		English (United States)			
Apache_error	Apache Error Log	Apache_error	.*\[(.*)\].*	EEE MMM dd HH:mm:ss yyyy		English (United States)			
Apache HTTP Server Access (for connector forwarder)	Apache HTTP Server Access log type used to be forwarded to streaming connector					English (United States)			
Apache HTTP Server Error (for connector forwarder)	Apache HTTP Server Error log type used to be forwarded to streaming connector					English (United States)			
audit_log	Syslog for Audit Log files	audit_log				English (United States)			
Bluecoat_proxy	Bluecoat Proxy SG	Bluecoat_proxy	.*(\w{3})s\d+\s\d\d\d\d	MMM dd HH:mm:ss		English (United States)			
CEF	for CEF event					English (United States)			
Cisco_PIX	Cisco PIX	Cisco_PIX	.*(\w{3})s\d+\s\d\d\d\d	MMM dd HH:mm:ss		English (United States)			

Figure 6-4 Source Types tab

The following source types have associated parsers:

Source type	Description
Apache_access	Apache Access Log
Apache_error	Apache Error Log
audit_log	Syslog for Audit Log files
Bluecoat_proxy	Bluecoat Proxy SG
Cisco_PIX	Cisco PIX
IBM_DB2	IBM DB2 9.x Audit Log
Juniper_NSM	Juniper NSM 2009 Syslog
logger_syslog	Syslog for syslog files on Logger appliance
Microsoft_DHCP	Microsoft DHCP for 2008 v6 log files
syslog	Simple Syslog
TippingPoint_SMS	Tipping Point SMS 2.5 Syslog
VMWare_ESX	VMWare ESX Syslog

Logger can forward an event to ESM by using a Connector forwarder, which then forwards it to a Streaming Connector. This connector normalizes the event and forwards it to ESM. If you need forward events to ESM by using a Connector forwarder, you must choose one of the following source types:

Source Type	
Apache HTTP Server Access	Juniper Steel-Belted Radius
Apache HTTP Server Error	Microsoft DHCP Log
IBM DB2 Audit	Other

To add a source type:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Source Types** tab (right panel).

The Source Types tab, shown in [Figure 6-4 on page 275](#), displays the current source types. You can sort the fields by clicking the column headers.

- 3 Click **Add**.
- 4 Fill in the fields to define the source type:

Table 6-8 Source Type Fields

Field	Description
Name	The name of the source type.
Description	A description of the source type.

Table 6-8 Source Type Fields (Continued)

Field	Description
Parser	The parser you want to associate with this source type. If the parser you need does not appear in the drop-down list, you can add one. For information on how to add a parser, see "Parsers" on page 278 .
Event Time Location	<p>A regular expression describing the timestamp in the log file. For example:</p> <pre>.*\[(.*?) \].*</pre> <p>This expression specifies that the timestamp is found inside the first set of square brackets on each line. The first capturing group (the part of the regex in parentheses) is the part that is then parsed using the Date/time format.</p> <p>You can specify that there is no timestamp in the log file with `.`.</p>
Event Time Format	<p>A regular expression describing the date and time format in the log file. For example, <code>dd/MMM/YYYY:HH:mm:ss Z</code></p> <p>You can specify that there is no timestamp in the log file with `.`.</p> <p>For more information about event time, see "Time Range" on page 86 and "Date and Time Specification" on page 273.</p>
Multiline Event Starts With	A regular expression describing how to recognize when adjacent lines are of the same event or when a new event starts. For example if each event starts with the date in the format, <code>yy-MM-dd HH:mm:ss.SSS</code> you could use <code>(\d+-\d+-\d+ \d+:\d+:\d+.\d+)</code> to indicate the start of a new event.
Locale	Select a locale from the pulldown list, such as English (United States), Chinese (Hong Kong), Chinese (Taiwan), and so on. This is locale of the data Logger should find in the file.


5 Click **Save**.

To edit a source type:


1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.

2 Click **Event Input** (left panel) > **Source Types** tab (right panel).

The Source Types tab, shown in [Figure 6-4 on page 275](#), displays the current source types. You can sort the fields by clicking the column headers.

3 Locate the source type that you want to update and click the Edit icon () on that row.



The Edit icon () is not available for out-of-box source types. You can copy the source type and make a similar one instead.

4 Edit the fields as appropriate.

Source type fields are documented in [Table 6-8 on page 276](#).


5 Click **Save**.

- 6 Disable and then re-enable any receivers that use this source type.




Changes in source type are not reflected in the associated receivers until you have re-enabled them.


To copy a source type:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Source Types** tab (right panel).
The Source Types tab, shown in [Figure 6-4 on page 275](#), displays the current source types. You can sort the fields by clicking the column headers.
- 3 Locate the source type that you want to copy and click the Copy icon () on that row.
- 4 Enter a name for the new source type and edit the fields as appropriate.
Source type fields are documented in [Table 6-8 on page 276](#).
- 5 Click **Save**.

To delete a source type:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Source Types** tab (right panel).
The Source Types tab, shown in [Figure 6-4 on page 275](#), displays the current source types. You can sort the fields by clicking the column headers.
- 3 Locate the source type that you want to delete and click the Remove icon () on that row.



The Remove icon () is not available for out-of-box source types. You can only remove source types that you added.

- 4 Click **OK** to confirm the removal.

Parsers

Parsers enable you to extract and manipulate raw events (non-CEF data) from different sources in your network environment. Once you have parsed event fields, you can easily search for data, chart it, and perform other operations on it. One user with in-depth knowledge of the events can create the parser, and then all users who look at those events will get the benefit of that work.

Parsers provide you with a simple way to read events. Instead of looking at raw event data and trying to figure out what it means, you can use a parser to extract portions of non-CEF events into fields. However, the fields created by the parser are available only for search operations, and are not added to the Logger schema.

You can use a parser either of the following ways:

- Use the parser with a source type—You can associate the parser with a source type to extract any set of fields in any kind of event. For more information, see [“Source Types” on page 274](#).
- Use the parse command in a search—During a search, you can use the parse command to extract fields from events and use other search operators (such as where, chart, top, etc) to further refine the search or manipulate the data in the fields. This is particularly useful for IT operations and other customers who need to extract and manipulate raw event data.

Using Parsers with Source Types

Logger provides a number of pre-configured parsers with associated source types. You can also define new parsers and associate them with source types. Only one parser can be associated with a source type, however, multiple source types can use the same parser. Out-of-box parsers cannot be edited or deleted, but you can copy them to make a similar parser to meet your needs. You can edit or delete custom parsers as desired.

Receivers	Source Types	Parsers	
Add			
Name	Parser Type	Description	Definition
Sample_Extract_Parser	Extract Parser	Sample Extract Parser	pair delimiter [\ :] key/value delimiter [=] fields [IN, MAC, DST,...
Apache_access	Rex Parser	Apache Access Log Parser	(?<SourceHost>\S+)[\s-]+(?<Identity>\S+)?[\s-]+\[(?<Date>...
Apache_error	Rex Parser	Apache Error Log Parser	\[(?<timestamp>\S+ \S+ \d{1,2} [\d:]+ \d{4}) \]\s+\[(?<severity>...
audit_log	Rex Parser	Syslog parser for Audit Log files	type=(?<type>\S+)\s+msg=audit\((?<epoch>\d+\.\d+):(?<e...
Bluecoat_proxy	Rex Parser	Bluecoat Proxy SG Parser	(?<Service>ProxySG): (?<ID>[A-Fa-f0-9]+) (?<Message>.*)...
Cisco_PIX	Rex Parser	Cisco PIX Parser	(?<priority>\d+)>)?(?<timestamp>\w{3}\s*\d+\d{0,4})\s?\S{...
IBM_DB2	Rex Parser	IBM DB2 9.x Audit Log Parser	(?<Timestamp>[\d\.\-]+)[^\w]?,[^\w]? (?<Category>\w*)[^\w]?,[^\w]...
Juniper_NSM	Rex Parser	Juniper NSM 2009 Syslog parser	(?<logDayId>\d+),\s+(?<logRecordId>\d+),\s+(?<nsmRece...
logger_syslog	Rex Parser	Syslog parser for syslog files on Logger appliance	(?<timestamp>\w{3} \d+ \d\d:\d\d:\d\d)? (?<host>[\w\d\.\-]+)?...
Microsoft_DHCP	Rex Parser	Microsoft DHCP parser for 2008 v6 log files	(?<EventId>\d+),(?<Date>\d+/\d+/\d+),(?<Time>\d+:\d+:\d+)...
syslog	Rex Parser	Simple Syslog Parser for syslog files on Linux	(?<priority>\d+)>)?(?<timestamp>\w{3}\s*\w{0,3}\s*\d+ \d{...
TippingPoint_SMS	Rex Parser	Tipping Point SMS 2.5 Syslog Parser	(?<timestamp>\w{3} \d+ \d\d:\d\d:\d\d)? (?<host>[\w\d\.\-]+)? ...
VMWare_ESX	Rex Parser	VMWare ESX Syslog Parser	(?<Module>\S+?vmware-hostd Hostd Fdm DCUI vmkwarn...

Figure 6-5 Parsers tab

Using the Parse Command

The `parse` command can be used to invoke a parser on any non-CEF events that are returned by a search. It applies the definition of the parser, such as the regular expression of a rex parser, to each event. Then it adds the fields that are extracted by that regular expression to the fields that are being passed through. For a REX parser, this is functionally the same as having a rex command with the same regular expression as the definition of the parser, so you can think of a REX parse command as invoking a saved rex expression.

For more information about the `parse` command, see [“parse” on page 541](#). For information about searching in general, see [“Searching and Analyzing Events” on page 75](#).

Working with Parsers

You can define two types of parsers—a REX parser or an Extract parser. Before adding the parser, you need to define the query you want to use for parsing events.

For a Rex parser, one way to do this is to use the rex search operator to test and adjust a regular expression until it returns the desired fields from the events that you want it to handle. Then copy the rex expression and paste it into the parser's Definition field. For an Extract parser, use the extract operator. For more information about the search operators, see ["parse" on page 541](#), ["rex" on page 545](#), and ["extract" on page 537](#).

The parser used in a search will be displayed in the Parser column of the search results. If the event was parsed, this field contains the name of the parser. If the event was not parsed successfully, this field contains "Not parsed". If no parser is defined for the source type or if there is no source type, the field is blank.

To add a parser:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Parsers** tab (right panel) or Click **Search** (left panel) > **Parsers** tab (right panel).

The Parsers tab, shown in [Figure 6-5 on page 279](#), displays the current parsers. You can sort the fields by clicking the column headers.

- 3 Click **Add**.
- 4 Enter a name for the parser.
- 5 Choose the Parser Type from the drop-down list.
- 6 Click **Save**.

The fields displayed in the Edit Parser dialog box according to the type of parser.

- 7 Fill in the fields for the parser.

Table 6-9 Parser fields

Field	Description
Name	The name of the parser. Enter a new name if you want to change the existing name.
Description	A meaningful description of the purpose of the parser.
Rex parsers only	
Definition	The rex expression that you want to use to parse events.
Extract parsers only	
Pair Delimiter	The characters separate key/value pairs within an event. Enter only the separator characters, for example: \\,
Key/Value Delimiter	The characters that separate the key from the value. Enter only the delimiter character, for example: =

Table 6-9 Parser fields (Continued)


Field	Description
Fields	<p>The list of field names to use when parsing events.</p> <p>Enter the field names, separated by comma (.). For example, to parse events like: <code>foo=abc, bar=xyz, baz=def</code></p> <p>Enter: <code>foo,bar,baz</code></p>

- 8 Click **Save**.


To edit a parser:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Parsers** tab (right panel) or
Click **Search** (left panel) > **Parsers** tab (right panel).

The Parsers tab, shown in [Figure 6-5 on page 279](#), displays the current parsers. You can sort the fields by clicking the column headers.

- 3 Locate the parser that you want to update and click the Edit icon () on that row.



The Edit icon () is not available for out-of-box parsers. You can copy the parser and make a similar one instead.

- 4 Edit the parser fields as appropriate.

The fields displayed in the Edit Parser dialog box according to the type of parser. Parser fields are documented in [Table 6-9 on page 280](#).

- 5 Click **Save**.

To copy a parser:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Parsers** tab (right panel) or
Click **Search** (left panel) > **Parsers** tab (right panel).

The Parsers tab, shown in [Figure 6-5 on page 279](#), displays the current parsers. You can sort the fields by clicking the column headers.

- 3 Locate the parser that you want to copy and click the Copy icon () on that row.

The fields displayed in the Edit Parser dialog box according to the type of parser.

- 4 Enter a name for the new parser and edit the fields as appropriate.

Parser fields are documented in [Table 6-9 on page 280](#).

- 5 Click **Save**.

To delete a parser:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Parsers** tab (right panel) or Click **Search** (left panel) > **Parsers** tab (right panel).

The Parsers tab, shown in [Figure 6-5 on page 279](#), displays the current parsers. You can sort the fields by clicking the column headers.
- 3 Locate the parser that you want to delete and click the Remove icon (✖) on that row.



The Remove icon (✖) is not available for out-of-box parsers. You can only remove parsers that you added.

- 4 Click **OK** to confirm the removal.

Example: Creating an Extract Parser

Suppose you want to create a parser to find the contents of the INT, MAC, DST, and SRC fields of a log like the one below.

```
Jul 12 14:30:31 n15-214-128-h92 kernel: IN=eth2
MAC=00:24:e8:60:cb:82:00:50:56:92:2a:d5:08:00 SRC=15.214.157.89 |
DST=15.214.128.92 LEN=52 TOS=0x00 PREC=0x00 TTL=128 ID=21408 DF
PROTO=TCP SPT=56978 DPT=443 WINDOW=8192 RES=0x00 SYN URGP=0

Jul 12 14:30:31 n15-214-128-h92 kernel: IN=eth2 |
MAC=00:24:e8:60:cb:82:00:50:56:92:2a:d5:08:00 | SRC=15.214.157.89 |
DST=15.214.128.92 LEN=52 TOS=0x00 PREC=0x00 TTL=128 ID=21408 DF
PROTO=TCP SPT=56978 DPT=443 WINDOW=8192 RES=0x00 SYN URGP=0

Jul 12 14:30:31 n15-214-128-h92 kernel: IN=eth2 |
MAC=00:24:e8:60:cb:82:00:50:56:92:2a:d5:08:00 | SRC=15.214.157.89 |
DST=15.214.128.92 LEN=52 TOS=0x00 PREC=0x00 TTL=128 ID=21408 DF
PROTO=TCP SPT=56978 DPT=443 WINDOW=8192 RES=0x00 SYN URGP=0
```

In this sample log, the field values are indicated with an equal sign (=), and fields are delimited by pipe (|) and colon (:). You could use the following query to search for the contents of the IN, MAC, DST, and SRC fields.

```
extract pairdelim= "|:" kvdelim= "=" fields= "IN,MAC,DST,SRC"
```

The following steps describe how to make an extract parser using that query.

To create an example extract parser:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Input** (left panel) > **Parsers** tab (right panel) or Click **Search** (left panel) > **Parsers** tab (right panel).
- 3 Click **Add**.

The Add Parser dialog box opens.

- 4 Enter a Name and select the Parser Type. For the example, enter:

Name: Sample_Extract_Parser

Parser Type: Extract Parser

- 5 Click **Save**.

The Edit parser dialog box opens.

- 6 Enter the Pair Delimiter, Key value, and Fields for the parser. For the example, enter:

Pair Delimiter: \ | \ :

Key/Value Delimiter: =

Fields: INT, MAC, DST, SRC



Note

You need to escape the pipe (|) and the colon (:) with backslash (\).

- 7 Click **Save**. The Parsers tab displays the new parser.

Receivers

Source Types

Parsers

Add

Name	Parser Type	Description	Definition			
Sample_Extract_Parser	Extract Parser	Sample Extract Parser	pair delimiter [:] key/value delimiter [=] fields [IN, MAC, DST,...			
Apache_access	Rex Parser	Apache Access Log Parser	(?<SourceHost>\S+)\S+(?<Identity>\S+)\S+ \[(?<Date>...			
Apache_error	Rex Parser	Apache Error Log Parser	\[(?<timestamp>\S+ \S+ \d{1,2} [d:]+ \d{4}) \S+\[(?<severity>...			
audit_log	Rex Parser	Syslog parser for Audit Log files	type=(?<type>\S+)\S+msg=audit\[(?<epoch>\d+\.\d+):(?<e...			
Bluecoat_proxy	Rex Parser	Bluecoat Proxy SG Parser	(?<Service>ProxySG): (?<ID>[A-Fa-f0-9]+) (?<Message>.*)...			
Cisco_PIX	Rex Parser	Cisco PIX Parser	(<(?<priority>\d+)>)?(?<timestamp>\w{3}\S*\d+ \d{0,4}\S?\S...			
IBM_DB2	Rex Parser	IBM DB2 9.x Audit Log Parser	(?<Timestamp>[\d\.\-]+)[^\w]?,[^\w]? (?<Category>\w*)[^\w]?,[^...			
Juniper_NSM	Rex Parser	Juniper NSM 2009 Syslog parser	(?<logDayId>\d+),\S+(?<logRecordId>\d+),\S+ (?<nsmRece...			
logger_syslog	Rex Parser	Syslog parser for syslog files on Logger appliance	(?<timestamp>\w{3} \d+ \d:\d:\d:\d:\d)? (?<host>[\w\d\.-]+)?(...			
Microsoft_DHCP	Rex Parser	Microsoft DHCP parser for 2008 v6 log files	(?<EventId>\d+),(?<Date>\d+/\d+/\d+),(?<Time>\d+:\d+:\d+)...			
syslog	Rex Parser	Simple Syslog Parser for syslog files on Linux	(<(?<priority>\d+)>)?(?<timestamp>\w{3}\S*\w{0,3}\S*\d+ \d{...			
TippingPoint_SMS	Rex Parser	Tipping Point SMS 2.5 Syslog Parser	(?<timestamp>\w{3} \d+ \d:\d:\d:\d:\d)? (?<host>[\w\d\.-]+)? ?...			
VMWare_ESX	Rex Parser	VMWare ESX Syslog Parser	(?<Module>\S+?vmware-host Hostd Fdm DCUI vmkwarn...			

Event Output

Use the Event Output section to manage the forwarders that send stored events to other destinations, including ArcSight Manager.

Forwarders

ESM Destinations

Certificates

Filter by Type

All

Add



















Name	Type	Filter Type	IP/Host	Port	Query				
esm-regex	ArcSight ESM (CEF) Forwarder	Regular Expression			Microsoft Snort Unix				
esm-unified	ArcSight ESM (CEF) Forwarder	Unified Query			deviceVendor CONTAINS "Microsoft" OR deviceVendor = "Unix"				
tcp1	TCP Forwarder	Unified Query	15.214.128.245	1768					
tcp2	TCP Forwarder	Regular Expression	15.214.128.245	1913					
udp-regex	UDP Forwarder	Regular Expression	15.214.128.245	1109	deviceGroup(FFR1,FFR2)				
unified-fwd	ArcSight ESM (CEF) Forwarder	Unified Query			deviceVendor = "Microsoft"				

Figure 6-6 Event Output screen

Forwarders

Forwarders send all events, or events which match a particular filter, on to a particular host. The ability to define a different filter for each forwarder allows Logger to divide traffic among several destinations. For example, because Logger can handle much higher event rates than ArcSight Manager, Logger might be used to forward events to a number of ArcSight Managers. Forwarder filters make it possible to split the flow between the Managers, using one forwarder for each Manager. Additionally, forwarding enables you to send a subset of events to other destinations for further processing while maintaining all events on Logger for long-term storage.

The forwarding filter is a query that searches for matching events, optionally within a time range. You can create two types of forwarder filters—continuous and time-range bound.

- A continuous filter constantly evaluates the incoming events and forwards the matching ones to the specified destination.
- A time-range bound filter uses a time range in addition to the specified condition to determine whether an event should be forwarded to the destination. If the event falls within the specified time range and matches the specified condition, it is forwarded; otherwise, it is not. The Logger receipt time of an event is used to determine whether an event will be forwarded to a destination when a forwarder filter specifies a time range by which events are evaluated for forwarding. Once a forwarder has forwarded all events within a time range, it does not forward any more events.

A forwarder only forwards events from the Logger that it is configured on; it cannot forward events from peer Loggers.

A forwarder's operation can be paused and resumed at any point in time. When a forwarder resumes operation, forwarding resumes from the last checkpoint that was established before the forwarding operation was paused.

You can also disable and re-enable a forwarder. When you re-enable a forwarder, all previously established checkpoints are removed and forwarding starts over again as per the forwarder configuration—forwarders with continuous filters start from the current time, while forwarders with time-range bound filters start from beginning of the configured time range.

Forwarder types include UDP Forwarder, TCP Forwarder, Connector Forwarder, and ArcSight ESM Forwarder:

- **UDP:** UDP forwarders forward events by using the User Datagram Protocol.
- **TCP:** TCP forwarders forward events by using the Transmission Control Protocol.
- **Connector Forwarder:** Connector forwarders send events to the Logger Streaming Connector.
- **ArcSight ESM:** ArcSight ESM forwarders send Common Event Format (CEF) events to an ESM Destination. The built-in connector on Logger is used to forward these events to ESM.

As a best practice, do not add more than 10 regular expression forwarders on a Logger. Even though each additional forwarder improves the forwarding rate, the relation is not proportional. In high EPS (events per second) situations or situations where other resource-intensive features are running in parallel (alerts, reports, and several search operations) and the forwarding filter is complex, adding too many forwarders may reduce performance because forwarders have to compete for the same Logger resources besides competing for the same built-in connector for forwarding.

Prior to Logger 5.2, you could only specify a regular expression query for the filter. However, starting with 5.2, you can also specify indexed search queries (known as Unified Queries). Doing so enables you to take advantage of the indexing technology to quickly and efficiently search for events to forward.



Unified query-based forwarders forward events once they have been indexed. Therefore, these forwarders can exhibit “bursty” behavior because indexing occurs in batches on Logger. You might notice the bursty behavior in the EPS out gauge (on top of the Logger interface screen)—the gauge will display high EPS level as a burst of data is forwarded and then drop back to normal level.

To create a forwarder:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 In the Forwarders tab, click **Add** to display the following form.

- 4 Enter a name for the new forwarder and choose the forwarder type appropriate for your need: UDP Forwarder, TCP Forwarder, Connector Forwarder, or ArcSight ESM (CEF) Forwarder type.
- 5 Select the type of forwarding filter you will specify for this forwarder—Unified or Regular Expression. Select “Unified” if you want to specify an indexed search query or “Regular Expression” to specify a regular expression query.
- 6 Click **Next**.
- 7 Enter additional, type-specific information as described in [Table 6-10, “Forwarder Parameters,” on page 287](#). Click **Save**.

- 8 New forwarders are initially disabled. Click the disabled icon (🚫) to enable the new forwarder.

Table 6-10 Forwarder Parameters

Parameter	Forwarder Types	Description
Name	All	The name that you entered in the previous screen is displayed automatically. If you want to change the name, make the change on this screen.
Query	All	<p>Enter the query that will be used to filter events that the forwarder will forward, or select a filter from the Filter list below.</p> <p>Forwarder queries can be constrained by device groups and storage groups, but not by Peers.</p> <p>If you selected Unified Query in the previous screen, enter an indexed search query that includes full-text and field-based indexed fields. You can click the Advanced Search link to access the Search Builder tool to build an indexed query. (See “Accessing the Advanced Search Builder” on page 97 for more information.)</p> <p>The unified query you specify must follow these guidelines, or you will not be able to save the query and thus the forwarder:</p> <ul style="list-style-type: none"> Queries in the following format are valid; all other formats are not allowed. <pre>(full-text terms field search)* regex</pre> <p>That is, the query must only contain full-text (keyword) and field-based query elements; it cannot contain any aggregation search operators, or operators that process the searched data further to refine the search. For example, chart, sort, eval, top, and so on.</p>
Query (Continued)	All	<p>Therefore, this is a valid query: failed message CONTAINS “failed device”</p> <p>However, this is an invalid query: failed message CONTAINS “failed device” sort deviceEventCategory</p> <p>The query can contain the <code>regex</code> operator after a pipeline character (<code> </code>). Therefore, this is a valid query for a forwarder: failed message CONTAINS “failed device” regex deviceEventCategory = “fan”</p> <ul style="list-style-type: none"> All search terms (except the “regex” portion) in a query must be indexed. If a query contains full-text (keyword) terms, full-text indexing must be enabled. Similarly, if the query contains a field, field-based indexing must be enabled and the specified field must be indexed. <p>If you selected Regular Expression in the previous screen, specify a regular expression in this text box. See “Searching for Events on Logger” on page 106.</p>



Table 6-10 Forwarder Parameters (Continued)

Parameter	Forwarder Types	Description
Filters	All	<p>Instead of specifying a unified query, you can select a filter from the Filters list. The Filters list contains all saved filters and the predefined system filters on your Logger. Select a filter that meets the validity guidelines described in “Query” on page 287. Otherwise, the user interface will display an error when you save the forwarder definition. <i>You can only select one unified query filter per forwarder.</i></p> <p>Similarly, when creating a regular expression based filter, select a filter from this list. <i>You can select multiple filters for a regular expression based forwarder.</i></p>
Filter by time range	All	<p>If you are creating a continuous filter, which continuously evaluates incoming events and forwards the matching ones, skip this parameter. In this case, the query is run continuously and forwarding continues until you pause it.</p> <p>If you are creating a time range bound filter, check this box to specify a time range of events that the forwarder will forward. If you enter a time range, the forwarder sends events that are within that time range and stops.</p> <p>When you check this box, the Start and End dates and Time fields are displayed.</p> <p>Start must be earlier than End. Specifying a time in the future changes that field to the current time. For example, specifying a Start of the current day at 7 a.m. and an End of current day at 7 p.m. will produce events with timestamps from 7 a.m. to the time the filter is saved (that is, earlier than 7 p.m.).</p>
Source Type	Connector	<p>Select from the pulldown list of log file types, including:</p> <ul style="list-style-type: none"> • Apache HTTP Server Access • Apache HTTP Server Error • IBM DB2 Audit • Juniper Steel-Belted Radius • Microsoft DHCP Log • Other <p>Note: The Source type must be the same in receiver, forwarder, and SmartConnector. See “Forwarding Log File Events to ESM” on page 295.</p> <p>A receiver can only receive events of a single source type. Set up separate receivers for each type of log file.</p>
Preserve Syslog Timestamp	UDP, TCP	<p>Set to true to preserve the syslog timestamp. The default is true--the timestamp is the original receipt time of the event.</p> <p>If set to false, original timestamp is replaced with Logger's receipt time.</p>
Preserve Original Syslog Sender	UDP, TCP	<p>Set to true to send the event as-is, without inserting Logger's IP address in the hostname (or equivalent) field of the syslog event. The default is true.</p> <p>If set to false, Logger's information is inserted in the hostname (or equivalent) field of the syslog event.</p>

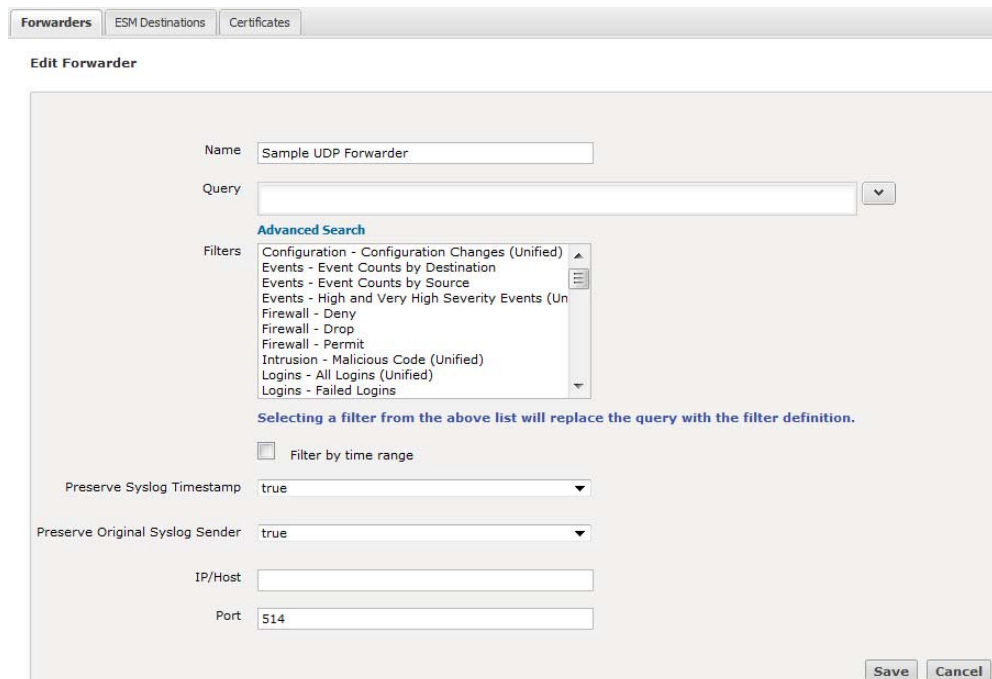
Table 6-10 Forwarder Parameters (Continued)

Parameter	Forwarder Types	Description
IP/Host	UDP, TCP, Connector	The IP address or host name of the destination that will receive forwarded events. Note: You cannot configure a Logger forwarder to send data to the same Logger on which it is configured.
Port	UDP, TCP, Connector	The port on the destination to which the forwarder will forward events. The default port is 514.
Connection Retry Timeout	TCP, Connector, ESM	The time, in seconds, to wait before retrying a connection. The default is 5 seconds.
ESM Destination	ESM	The ESM Destination for the target Manager. (See “ESM Destinations” on page 291.)

To edit a forwarder:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 In the Forwarders tab, locate the forwarder you want to edit.
- 4 If the forwarder is enabled, click the enabled () icon to disable it.
- 5 Click the Edit icon ().

The following screen shows the Edit Forwarder screen for a regular expression based forwarder. The Edit Forwarder screen for a Unified Query forwarder lists the Unified Query based filters and the Query text box only allows you to specify one query.



Forwarders ESM Destinations Certificates

Edit Forwarder

Name: Sample UDP Forwarder

Query:

Advanced Search

Filters:

- Configuration - Configuration Changes (Unified)
- Events - Event Counts by Destination
- Events - Event Counts by Source
- Events - High and Very High Severity Events (Un
- Firewall - Deny
- Firewall - Drop
- Firewall - Permit
- Intrusion - Malicious Code (Unified)
- Logins - All Logins (Unified)
- Logins - Failed Logins

Selecting a filter from the above list will replace the query with the filter definition.

☐ Filter by time range


Preserve Syslog Timestamp: true

Preserve Original Syslog Sender: true



IP/Host:

Port: 514


Save Cancel

- 6 Edit the information in the form, as described in [Table 6-10 on page 287](#), and click **Save**.
- 7 Click the disabled icon () to re-enable the forwarder and commit the changes.


To delete a forwarder:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 Locate the forwarder that you want to delete in the table.
- 4 If the forwarder is enabled, click the enabled () icon to disable it.
- 5 Click the Delete icon ()
- 6 Click **OK** to confirm the delete.


To pause a forwarder:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 Locate the forwarder that you want to pause from the list of forwarders.
- 4 Click the Pause icon ()

To resume a forwarder:


- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 Locate the forwarder whose operation you want to resume.
- 4 Click the Resume icon ()

To disable a forwarder:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 Locate the forwarder that you want to disable.
- 4 Click the enabled icon ()

To enable or re-enable a forwarder:

Wait a few minutes to disable a forwarder that was just enabled. Likewise, wait before enabling a forwarder that has just been disabled. Background tasks initiated by enabling or disabling a forwarder can produce unexpected results if they are interrupted.

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Output** in the left panel.
- 3 Locate the forwarder that you want to enable or re-enable.
- 4 Click the disabled icon ()

ESM Destinations

An ESM Destination establishes a connection to an ArcSight Manager so that you can forward events (and alerts) from the Logger to the Manager using Logger's built-in SmartConnector. The SmartConnector sends CEF events, which are already normalized or categorized.

For more information about the Common Event Format (CEF), refer to Implementing ArcSight CEF. For a downloadable a copy of this guide, search for "ArcSight Common Event Format (CEF) Guide" on the Protect 724 Community at <https://protect724.arcsight.com>.

Logger can forward these types of events to an ArcSight Manager:

- Syslog events to an ArcSight Syslog SmartConnector that is connected to an ArcSight Manager
- Common Event Format (CEF) events directly to an ArcSight Manager using Logger ESM Destinations. An ESM Destination appears as a SmartConnector to an ArcSight Console.
- Events received by file receivers where the type specified is not Other. Such events are forwarded using the ArcSight Streaming SmartConnector.

Maximum number of ESM Destinations that can be configured: As many allowable on the SmartConnectors you are using.



Note

Do not use basic aggregation for Logger's built-in SmartConnector because it is resource intensive. (Basic aggregation is set using the Enable Aggregation (in seconds) field from the ArcSight Console.) Instead, follow these steps on the ArcSight Console to configure field-based aggregation:

- 1 Ensure that Processor > Enable Aggregation (in seconds) is set to "Disabled" (to disable basic aggregation).
- 2 Right-click the connector and select **inspect/edit/**.

For additional details about configuring field-based aggregation, refer to the ArcSight SmartConnector's User's Guide.

To setup Logger to forward events to an ArcSight Manager:

- 1 Copy the server SSL certificate file from an ArcSight Console or other component that is already communicating with the target Manager, and upload the certificate file to Logger, as described ["Uploading a Certificate to the Logger:" on page 294](#).

If your Logger operates in FIPS mode, a valid and current (non-expired) server SSL certificate file from the ArcSight Manager is required on the Logger; otherwise, the forwarder will not forward events to it.



Note

You cannot import the cacerts file, which is a repository of trusted certificates, to the Logger. Instead, you need to import specific SSL certificate files.

- 2 Create an ESM Destination, as described in ["To create an ESM Destination:" on page 292](#).
- 3 Create an ESM forwarder that refers to this ESM Destination. (See ["Forwarders" on page 285](#)).

Figure 6-7 ESM Destinations page

To create an ESM Destination:

Make sure you have loaded the certificate file for ArcSight Manager as described in [“Uploading a Certificate to the Logger:” on page 294](#) before adding it as a destination on the Logger. If the certificate file does not exist on the Logger, you will not be able to create an ESM Destination.

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Output** in the left panel.

Click **Alerts** in the left panel if you are creating an ESM Destination for forwarding Alerts.



Note

The ESM Destinations tab located under Event Output and Alerts in the left panel is the same and contains all ESM Destinations configured on a Logger. The tab is accessible from two UI locations to ease configuration.

- 3 In the **ESM Destinations** tab, click **Add**. The page shown in [Figure 6-7](#) is displayed.
- 4 Enter the following parameters:

Parameter	Description
Name	The name for this ESM Destination.
Connector Name	The SmartConnector name.
Connector Location	The physical location of the SmartConnector machine. If you do not want to specify a location, enter “None.”
Logger Location	The physical location of the Logger. If you do not want to specify a location, enter “None.”

Parameter	Description
IP or Host	The ArcSight Manager to which the forwarder will direct events. Note: Make sure the name or IP address you specify in this field is exactly the name or IP address configured on the ArcSight Manager. If the two names or IP addresses do not match, you will not be able to set up an ESM Destination successfully.
Port	Typically 8443.
User Name	The name of an existing User of the ArcSight Manager with administrator privileges.
Password	The password for the Login user. This password cannot contain the special characters percent (%), equal to (=), semicolon (;), double quote ("), single quote ('), less than (<), or greater than (>). While ArcSight Manager allows these special characters in passwords, Logger does not. If the ArcSight Manager user's password contains those characters, you will need to change the password in ArcSight Manager before configuring this password.

5 Click **Save**.



If you receive the following error when adding a new ESM Destination, make sure the host name you specified in the IP or Host field exactly matches the name configured on the ArcSight Manager. "

There was a problem: Failed to add destination

Additionally, if the ArcSight Manager is configured using a host name instead of IP address, make sure you add the ArcSight Manager host name and IP address in the Logger's hosts file (**System Admin > Network > Hosts**).

To delete an ESM Destination:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Event Output** in the left panel.

Click **Alerts** in the left panel if you are deleting an ESM Destination for forwarding Alerts.



The ESM Destinations tab located under Event Output and Alerts in the left panel is the same and contains all ESM Destinations configured on a Logger. The tab is accessible from two UI locations to ease configuration.

- 3 In the **ESM Destinations** tab, locate the ESM Destination that you want to delete and click the delete icon (✖) on that row.
- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the ESM Destination.

Certificates

Uploading a Certificate to the Logger:

You need to upload a valid server SSL certificate file for the ArcSight Manager that you are establishing as a Logger destination for forwarding events and alerts.

If your Manager does not have FIPS 140-2 mode enabled, you can obtain a certificate file for your Manager in these ways:

- From the Manager's keystore
- From the ArcSight Console's truststore
- From the truststore of one of the SmartConnectors that communicates with the Manager

Use the `keytoolgui` utility to export a Manager's certificate as described in the "Using Keytoolgui to Export Certificate" procedure in the ArcSight ESM Administrator's Guide. For detailed information about keystore, truststore, their locations on the Manager, ArcSight Console, and the SmartConnectors, see the ArcSight ESM Administrator's Guide. Once you have exported a certificate for your Manager, copy it to the machine from which you connect to your Logger.

If your Manager has FIPS 140-2 mode enabled, run this command to export the Manager's certificate from the Manager's `<ARCSIGHT_HOME>/bin` directory:

```
arcsight runcertutil -L -n managerkey -r -d  
<ARCSIGHT_HOME>/config/jetty/nssdb -o  
<absolute_path_to_manager.cert>
```

This command generates the `manager.cert` file, the Manager's certificate, in the location that you specified in the above command.



Note

By default, the `manager.cert` file will be exported to your `<ARCSIGHT_HOME>` directory if you do not specify the absolute path to `manager.cert` file destination.

To upload a certificate file for an ESM Destination:

- 1** Make sure you have copied the Manager certificate to the machine from which you connect to your Logger.
- 2** Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 3** Click **Event Output** in the left panel. Then click **Certificates**.

If you are creating an ESM Destination for forwarding Alerts, click **Alerts** in the left panel. Then click **Certificates**.
- 4** In the **Certificates** tab, click **Add** to display the following screen.

- 5 Enter an alias for the certificate file. This name is used to easily identify a certificate file. For example, `arcsight_esm_manager1_cert`.
- 6 Click **Browse** to locate the Manager certificate file you copied.
- 7 Check the “Overwrite Certificate” box if you want this certificate to overwrite an existing certificate with the same alias.
- 8 Click **Save**.

Forwarding Log File Events to ESM

Logger can read events from a log file and forward those events to a Logger Streaming SmartConnector that sends the events on to ArcSight Manager.

To forward log file events to ESM, configure the receiver, forwarder, and SmartConnector to accept the same source type (as described in [“Date and Time Specification” on page 273](#)).

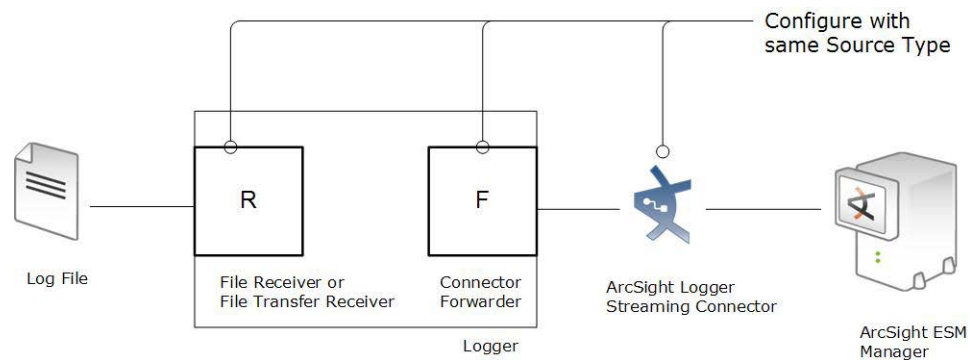


Figure 6-8 Configure the receiver, forwarder, and SmartConnector with the same source type to use Logger to forward log file events to ArcSight ESM.

Unlike events that Logger receives, such as syslog, SmartMessage, or CEF, log file events must be parsed to determine event timestamp. Therefore, if you need forward events to ESM by using a Connector forwarder, you must choose one of the following source types for the receiver:

Source Type

Apache HTTP Server Access	Microsoft DHCP Log
Apache HTTP Server Error	Other

Source Type	
IBM DB2 9.x Audit Log	Tipping Point SMS 2.5 Syslog
IBM DB2 Audit	VMWare ESX Syslog
Juniper Steel-Belted Radius	

Alerts

Alerts respond to events or specified event patterns with optional notification. Event patterns are specified events that occur above a particular frequency (a threshold number of events in a specified time period).

Alerts can be generated for internal events such as storage capacity warnings or, on some Logger appliance models, CPU temperature warnings, or for user-determined event patterns such as an alert is generated when five events from a specific device contain the word "unauthorized" within a five minute interval.

Logger provides two types of alerts:

- Real time alerts, discussed in ["Configuring and Managing Real Time Alerts" on page 299](#).
- Saved Search Alerts, discussed in ["Creating and Managing Saved Search Alerts" on page 302](#).

The following table compares the two types of alerts.

Real Time Alerts	Saved Search Alerts
No limit on the number of alerts that are defined. A maximum of five alerts can be enabled at any time.	Any number of alerts can be defined. All defined alerts are enabled and effective, however, a maximum of 50 alerts can run concurrently.
No limit on the number of configured e-mail destinations; however, you can only set one SNMP, one Syslog, and one ESM Destination.	No limit on the number of configured e-mail destinations; however, you can only set one SNMP, one Syslog, and one ESM Destination.
Only regular expression queries can be specified for these alerts.	Queries for these alerts are defined using the flow-based search language that allows you to specify multiple search commands in a pipeline format, including regular expressions. Aggregation operators such as chart and top cannot be included in the search query.
Alerts are triggered in real time. That is, when specified number of matches occurs within the specified threshold, an alert is immediately triggered.	These alerts are triggered at scheduled intervals. That is, when a specified number of matches occurs within the specified threshold, an alert is triggered at the next scheduled time interval .

Real Time Alerts	Saved Search Alerts
<p>To define a real time alert, you specify a query, match count, threshold, and one or more destinations.</p> <p>A time range is not associated with the queries defined for these alerts. Therefore, whenever the specified number of matches occurs within the specified threshold, an alert is triggered.</p>	<p>To define a Saved Search Alert, you specify a Saved Search (which is a query with a time range), match count, threshold, and one or more destinations.</p> <p>A time range (within which events should be searched) is specified for the query associated with these alerts. Therefore, specified number of matches within the specified threshold (in minutes) must occur within the specified time range. You can also use dynamic time range (for example, \$Now-1d, \$Now, and so on).</p> <p>For example, if a Saved Search query has these start and end times:</p> <p>Start Time: 5/11/2010 10:38:04</p> <p>End Time: 5/12/2010 10:38:04</p> <p>And, the number of matches and threshold are the following:</p> <p>Match count: 5</p> <p>Threshold: 3600</p> <p>Then, 5 events should occur in one hour anytime between May 11th, 2010 10:38:04 a.m. and May 12th, 2010 10:38:04 for this alert to be triggered.</p>

Alert Triggers and Notifications

An alert is triggered if a specified number of matches occurs within the specified threshold (time interval in seconds). When an alert is triggered, Logger creates an alert event containing the triggering events or event IDs, and sends notification through previously configured destinations—e-mail addresses, SNMP server, Syslog server, and ArcSight Manager.

By default, only alert notifications sent to e-mail destinations include all matching events that triggered the alert. You can configure your Logger to include matched events for SNMP, Syslog, and ESM Destinations as well. However, that kind of configuration is only possible through the command-line interface of the Logger; therefore, please contact customer support for instructions.

When are Alert events triggered?

You also specify a time window and a number of matching events. When that number of matching events is detected within the time window, an alert event is triggered.

Logger resets the count after detecting 100 matching events. Therefore, all events that occur in the time window will not necessarily be recorded in an alert. For example, if you configure the alert to be sent when there are 20 matching events in 2 minutes, and 152 events occur within two minutes, you will get 7 alerts, and 12 matching events will not be included in any alert. In this situation, the following alert events are triggered:

- Alert 1 has 20 matching events.
- Alert 2 has 40 matching events.
- Alert 3 has 60 matching events.
- Alert 4 has 80 matching events.
- Alert 5 has 100 matching events (1-100).

- Alert 6 has 20 matching events (101-120).
- Alert 7 has 40 matching events (101-140).

The remaining 12 events are being held, waiting to meet the threshold of 20 more events in a 2 minute interval.

Receiving Alert Notifications

In order to receive notification of an alert, set up the alert to be sent to a previously configured destination, such as an e-mail address, SNMP server, Syslog server, and ArcSight Manager.

By default, only alerts to e-mail destinations include all matched events that triggered the alert. You can configure your Logger to include matched events for SNMP, Syslog, and ESM Destinations as well. However, such a configuration is only possible through the command-line interface of the Logger; therefore, please contact customer support for instructions.

Sending Notifications to E-mail Destinations

When you send notifications for an alert via e-mail, the e-mail message contains both the trigger alert information and the matched (base) events.

The following is an example of the trigger alert information:

```
Alert event match count [1], threshold [10] sec
```

And the matched event:

```
Event Time [Tue Nov 11 16:46:49 PST 2008]
```

```
Event Receipt Time [Tue Nov 11 16:46:50 PST 2008]
```

```
Event Device Address [192.168.35.50]
```

```
Event Content [Dec 11 10:31:20 localhost
CEF:0|NetScreen|Firewall/VPN||traffic:1|Permit|Low| eventId=590
msg=start_time\= "2004-07-28 15:25:02" duration\=15 policy_id\=0
service\=SSH proto\=6 src zone\=Trust dst zone\=Untrust
action\=Permit sent\=656 rcvd\=680 src\=10.0.111.46
dst\=10.0.113.50 src_port\=54759 dst_port\=22 translated
ip\=192.91.254.2 port\=54759 app=SSH proto=TCP in=680 out=656
categorySignificance=/Normal categoryBehavior=/Access
categoryDeviceGroup=/Firewall categoryOutcome=/Success
categoryObject=/Host/Application/Service art=1165861874880
cat=Traffic Log deviceSeverity=notification act=Permit
rt=1165861874880 shost=n111-h046.qa.arcsight.com src=10.0.111.46
sourceZoneURI=/All Zones/System Zones/Private Address
Space/RFC1918: 10.0.0.0-10.255.255.255
sourceTranslatedAddress=192.91.254.2 sourceTranslatedZoneURI=/All
Zones/System Zones/Public Address Space/192.0.3.0-192.167.255.255
spt=54759 sourceTranslatedPort=54759 dst=10.0.113.50
destinationZoneURI=/All Zones/System Zones/Private Address
Space/RFC1918: 10.0.0.0-10.255.255.255 dp]
```

Sending Notifications to Syslog and SNMP Destinations

When configuring Logger to send alerts to SNMP and Syslog destinations, you should be familiar with this information:

- Logger supports SNMP 2.0.
- Unlike an e-mail alert, a trigger alert is sent separately from the alert that contains the matched (base) events that triggered the alert.
- All SNMP alerts are sent as SNMP traps; therefore, trigger alerts and their associated matched (base) events are received as SNMP traps on an SNMP destination. The SNMP trap includes the trigger event, but it does not include the events that caused the alert to trigger (matched events). The trigger event does include the event IDs of all the matched events. You can use the event IDs in the trigger alert to identify the associated matched events.

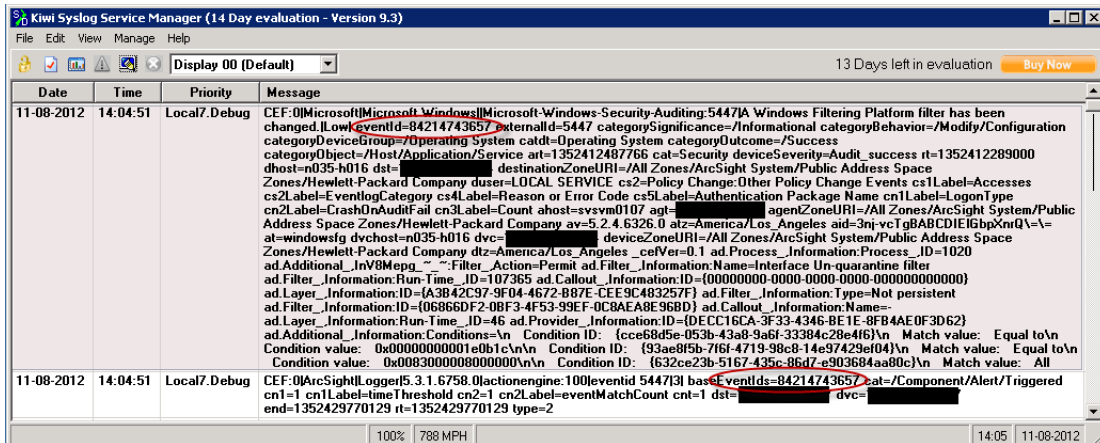


Figure 6-9 A triggered alert event and matching base event shown in Kiwi Syslog Service Manager



Non-CEF events do not contain event IDs. If you need to associate such base events with their trigger alert, send such events to Logger through a connector.

- SNMP uses UDP to send packets. As a result, the order in which alerts arrive at an SNMP destination is not guaranteed.
- When Syslog events are sent using UDP, the order in which the trigger alert and matched events arrive is not guaranteed.

Configuring and Managing Real Time Alerts

This section describes ways to configure and manage real time alerts. For information on Saved Search Alerts, see [“Creating and Managing Saved Search Alerts” on page 302](#).

Creating a Real Time Alert

To create an Alert, you will need to specify a query or filter, event aggregation values (Match count and Threshold, described below), and (optional) one or more notification destinations. If the new Alert will send notification using an SNMP, Syslog, or ESM Destination, set up those destinations before creating the Alert. To configure the e-mail destination, see [“Static Routes” on page 361](#). See also [“Sending Notifications to SNMP Destinations” on page 307](#), [“Sending Notifications to Syslog Destinations” on page 308](#), and [“ESM Destinations” on page 291](#).

When you create an alert, it is in disabled state. You can enable it using instructions in [“To Enable or Disable a Real Time Alert:” on page 301](#).

To create a Real Time Alert:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click **Add**. The page shown in [Figure 6-10 on page 301](#) is displayed.
- 4 Enter a name for the new Alert, specify a query or select an available Filter from the list. Events that match this query are candidates for the Alert. Alphanumeric characters and spaces are acceptable, however, some special characters such as % and & are not. For more information on Filters, see ["Filters" on page 312](#).

You can only specify regular expression queries for real time Alerts. However, a query expression for a scheduled saved alert can contain multiple search commands in a pipeline format, including regular expressions. Aggregation operators such as chart and top cannot be included in the search query. For more information about specifying a regular expression query, see ["The Need to Search Events" on page 75](#).



To test the validity of an alert query, use the Search user interface. Enter the query in the Search text box in the following format:

Real time Alert: |regex "regex expression"

Scheduled saved alert: _deviceGroup IN ["192.168.22.120 [TCPC]"]
name="*[4924TestAlert]*" AND ("192.168.*" OR categoryBehavior
CONTAINS Stop)

If the query is valid, cut and paste the regular expression between the double quotes (" ") in the Query text box on the Add Alert page.

- 5 Enter Match count and Threshold values. If the number of candidate events equals or exceeds the Match count within the Threshold number of seconds, the Alert will be triggered.

If you want to be notified when any event matches the filter (for example, for an internal event such as High CPU Temperature), enter a Match count of 1 and a Threshold of 1.



To maintain an optimal size of an alert event, the event does not contain event IDs of all the triggering events if you specify **Match count of 101 or higher**. As a result, the `baseEventCount` field in the event does not reflect the true number of matching events for such alert events.



Triggering events are truncated in multiples of 100. Therefore, if you specify a Match count of 101, only one event is included in the alert event and the `baseEventCount` field value is 1. Similarly, if you specify a Match count of 720, only 20 events are included and the `baseEventCount` field value is 20.

- 6 Enter notification destinations. Enter any combination of:
 - ◆ One or more e-mail addresses, separated by commas
 - ◆ An SNMP Destination—for more information, see ["Sending Notifications to SNMP Destinations" on page 307](#).
 - ◆ A Syslog Destination—for more information, see ["Sending Notifications to Syslog Destinations" on page 308](#).
 - ◆ An ArcSight Manager—for more information, see ["Sending Notifications to ESM Destinations" on page 309](#).
- 7 Click **Save**.

Figure 6-10 Add Alert dialog

To Enable or Disable a Real Time Alert:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Alerts** in the left panel.


Locate the Alert that you want to disable or enable. Click the associated icon ( or ) to enable or disable the Alert.



Note

A maximum of five alerts can be enabled at one time. To enable an additional alert, you will need to disable a currently enabled alert.

To Edit a Real Time Alert:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Locate the Alert that you want to edit and click the Edit icon () on that row.

A screen similar to that shown in [Figure 6-10 on page 301](#) appears. Remember that only alphanumeric characters can be used in an Alert name.

To Remove a Real Time Alert:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Alerts** in the left panel.

- 3 Locate the Alert that you want to remove and click the remove icon (✖) on that row.
- 4 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Alert.

To view Real Time Alerts:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.

Click **Alerts** in the left panel. The Alerts list is displayed, as shown in Figure 6-11.










Alerts						
Add						
If you are looking for scheduled alerts, you can find them on the Scheduled Searches/Alerts page.						
Name	Email Destination(s)	SNMP Destination	Syslog Destination	ESM Destination	Query	
IT Governance - Access Right Removed		NONE	NONE	NONE	:NOT: storageGroup(Internal Event Storage Group) :AND: categoryOutcome=/Success :AND: categoryBehavior= (/Authentication/Delete/Authorizati...	 ✖ ⌵
IT Governance - Access Right Removed [Import]		NONE	NONE	NONE	:NOT: storageGroup(Internal Event Storage Group) :AND: categoryOutcome=/Success :AND: categoryBehavior= (/Authentication/Delete/Authorizati...	 ✖ ⌵
IT Governance - Account Lockout		NONE	NONE	NONE	:NOT: storageGroup(Internal Event Storage Group) :AND: ^CEF:d{1}\Microsoft\Microsoft Windows\.*?Security: (539)644\}	 ✖ ⌵
IT Governance - Audit Log Cleared		NONE	NONE	NONE	:NOT: storageGroup(Internal Event Storage Group) :AND: ^CEF:d{1}\Microsoft\Microsoft Windows\.*? Security:517\}	 ✖ ⌵
IT Governance - Default Vendor Account Used		NONE	NONE	NONE	:NOT: storageGroup(Internal Event Storage Group) :AND: (d)suser=(admin)root isa [nobody guest manager sys system oracle orcladmin cisco...	 ✖ ⌵
IT Governance - Disallowed Port Access		NONE	NONE	NONE	:NOT: storageGroup(Internal Event Storage Group) :AND: :NOT: (d)spt=(80 443) :AND: dpt=d{1,} :AND: ^CEF:d{1}	 ✖ ⌵
IT Governance - Exploit of Vulnerability Detected		NONE	NONE	NONE	:NOT: storageGroup(Internal Event Storage Group) :AND: categoryTechnique=/Exploit/Vulnerability	 ✖ ⌵
IT Governance - Failed File Access		NONE	NONE	NONE	:NOT: storageGroup(Internal Event Storage Group) :AND: categoryOutcome=/Failure :AND: categoryObject=/Host /Resource/File	 ✖ ⌵
IT Governance - Failed File Deletion		NONE	NONE	NONE	:NOT: storageGroup(Internal Event Storage Group) :AND: categoryOutcome=/Failure :AND: categoryObject=/Host /Resource/File :AND: categoryBe...	 ✖ ⌵

Figure 6-11 Alert list

Creating and Managing Saved Search Alerts

This section describes ways to configure and manage Saved Search alerts. For information on real time alerts, see [“Configuring and Managing Real Time Alerts” on page 299](#).

Saved Search Alerts are based on the search queries that you have saved on Logger. For detailed information about Saved Search queries, see [“Saved Searches” on page 315](#). For each Saved Search Alert, you configure a match count, threshold, destination, and a schedule at which the alert will be triggered (if specified number of matches occurs within the specified threshold).



Note

To ensure system performance, a maximum of 200 alerts are allowed per saved search alert job. Therefore, if a saved search alert job triggers more than 200 alerts, only the first 200 alerts are sent out for that job iteration; the rest are not sent. Additionally, the job is aborted so it does not trigger more alerts for that iteration and the status for that job is marked “Failed” in the Finished Tasks tab (**Configuration** (or **Configuration > Settings**) > **Scheduled Tasks > Finished Tasks**). The job runs as scheduled at the next scheduled interval and alerts are sent out until the maximum limit is reached.


This limit does not exist on the real-time alerts.

Creating a Saved Search Alert

You can create a Saved Search Alert in two ways:

- From the search results page (**Analyze > Search**)
- From the Scheduled Searches page (**Configuration** (or **Configuration > Settings**) > **Saved Search > Scheduled Searches**)

To create a Saved Search Alert from the search results page:

- 1 Run a search, as described in [“Searching for Events on Logger” on page 106](#).
- 2 Click the Save icon () and enter the following settings.

Parameter	Description
Name	A name for the query you are saving.
Save as	Whether to save the query as a filter or as a Saved Search. To save the query as a Saved Search Alert, select “Saved search”.
Schedule it	Whether to schedule the alert right now or later. Click to schedule now, or leave blank to schedule later.
Schedule type	Whether the query is being saved as a scheduled search or as a scheduled alert. Scheduled searches run on a predetermined schedule and export results to a pre-specified location. Scheduled alerts run a search on a predetermined schedule and generate an alert if the specified number of events within the specified threshold is found.
Overwrite	If a query with the same name exists, whether that query should be overwritten. If you check this setting and a query with the same name exists, the existing query is overwritten with the one you are currently saving. If you do not check this setting, a warning message is displayed that prompts you to enter another name for the query.

- 3 Click **Save**.

If you checked the “Schedule it” setting in the previous step, you are prompted to choose if you want to edit the schedule, as follows. If you click **OK**, the Edit Scheduled Search page is displayed, as shown in the next step. If you click **Cancel**, the search is saved as a Saved Search but it is not scheduled to run.

- 4 If you checked the Schedule it setting previously, the Edit Scheduled Search page is displayed. This page enables you to define a schedule for the Saved Search job and alert options.

[Saved Searches](#)
[Scheduled Searches/Alerts](#)
[Saved Search Files \(Logger\)](#)

Edit Scheduled Search

Name

Schedule

Hour of day Hours (24 hour format)

Saved Searches

SaveIt
 SL_Saved_Search
 SL_SavedSearch

Use ctrl-click to select or deselect items

Job type

Search Result Export Options

Export Options
 ☐ Export to remote location
 ☒ Save to Logger

File format

Export directory name

Fields

Event Time, Receipt Time, Device, Logger, Name, Version, Device Vendor, Device Product, Device Version, Signature ID, Severity

☒ All fields

Include summary ☐

Include only CEF events ☐

To create a Saved Search Alert from the Scheduled Searches page:

- 1** Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2** Click **Saved Search** in the left panel.
- 3** Click **Scheduled Searches** in the right panel.
- 4** Click **Add**.

Saved Searches

Scheduled Searches/Alerts

Saved Search Files (Logger)

Edit Scheduled Search

Name

SaveIt job

Schedule

Everyday

Hour of day

3

Hours (24 hour format)

Saved Searches

SaveIt

SL_Saved_Search

SL_SavedSearch

Job type

Alert

Alert Options

Match count

Threshold (sec)

Email address(es)

SNMP destination

NONE

Syslog destination

NONE

ESM destination

NONE

Save

Cancel

5 Enter the following information.

Parameter	Description
Name	A name for the Saved Search you are saving.
Schedule	<p>Choose Everyday or Days of Week from the first pulldown menu.</p> <p>If Everyday:</p> <ul style="list-style-type: none">• EITHER select Hour of Day to specify the hour of the day in 24-hour format• OR select Every to specify the number of hours or minutes after which a Saved Search is performed. Select from the pulldown on the right side of Every to specify Hours or Minutes. By default, the number of hours and number of minutes is set to 15. <p>If Days of Week, enter the days (day 1 is Sunday) in the text box. Then choose Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the second text box.</p> <p>For example, to perform the Saved Search every day at 2 a.m., select Everyday in the first pulldown menu, then choose Hour of Day from the second pulldown menu and enter 2 in the text box. To perform the Saved Search every day at 2 a.m. and 3 p.m., enter 2,15 in the text box.</p> <p>For another example, to perform the Saved Search Tuesdays and Thursdays at 10 p.m., select Days of Week from the first pulldown menu and enter 3,5 for days. Then choose Hour of Day from the second pulldown menu and enter 22 in the text box.</p>

Parameter	Description
Saved Searches	Select from the list of Saved Searches. If a Saved Search to suit your needs does not exist in the list, click the Saved Searches tab (to the left of Scheduled Searches tab) to define it. For more information about defining a Saved Search query, see “Saved Searches” on page 315 . Note: You can only select one Saved Search for each Alert you configure.
Job Type	Select Alert for a Saved Search Alert.
Alert Options	
Match count	Number of events that should be matched in Threshold number of seconds for an alert to be triggered.
Threshold (sec)	Number of seconds within which the “Match count” events should be matched for an alert to be triggered.
Notification destinations are optional. If none is specified, a notification is not sent.	
Email address(es)	(Optional) A comma-separated list of email addresses to which the alert will be sent
SNMP destination	(Optional) An SNMP destination to which the alert will be sent. For more information, see “Sending Notifications to SNMP Destinations” on page 307 .
Syslog destination	(Optional) A syslog server address to which the alert will be sent. For more information, see “Sending Notifications to Syslog Destinations” on page 308 .
ESM Destination	(Optional) An ArcSight Manager address to which the alert will be sent. An ArcSight Manager—for more information, see “Sending Notifications to ESM Destinations” on page 309 .

6 Click **Save**.

7 Once a Saved Search Alert is created, you need to enable it. See [“To Enable or Disable a Saved Search Alert:” on page 306](#).



To Enable or Disable a Saved Search Alert:

1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.

2 Click **Saved Search** in the left panel.

3 Click **Scheduled Searches** in the right panel.

4 Identify the Saved Search Alert that you want to enable.

5 Click the associated icon ( or ) to enable or disable the alert.


To edit a Saved Search Alert:

1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.

2 Click **Saved Search** in the left panel.


3 Click **Scheduled Searches** in the right panel.

4 Locate the Saved Search Alert that you want to edit.

5 Click the Edit icon () and edit the information. For details about the settings, see [“To create a Saved Search Alert from the Scheduled Searches page:” on page 304](#).

- 6 Click **Save**.

To remove a Saved Search Alert:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.
- 4 Identify the Saved Search Alert that you want to remove.
- 5 Click the remove icon ().
- 6 Click **OK** to confirm the removal, or click **Cancel** to keep the alert.

To view Saved Search Alerts:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.

A list of the currently configured Saved Search Alerts is displayed.

Sending Notifications to SNMP Destinations

SNMP Destinations describe how Alert notifications should be sent using Simple Network Management Protocol (SNMP). Set up SNMP Destinations before creating Alerts that will use them. Before configuring an SNMP destinations, you should be familiar with the information in [“Sending Notifications to Syslog and SNMP Destinations” on page 298](#).

To Add an SNMP Destination:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **SNMP Destinations** tab in the right panel.
- 4 Click the **Add** button.
- 5 Enter parameters:

Parameter	Description
SNMP Destination Name	A name for this destination.
Connector Name	The SmartConnector name.
Connector Location	The physical location of the SmartConnector machine. If you do not want to specify a location, enter “None”.
Logger Location	Optional comment describing Logger’s physical location.
SNMP Host	Host name or IP address.
SNMP Port	162, by default.
Community Name	SNMP community name.

- 6 Click **Save** to create the new SNMP Destination.

To Remove an SNMP Destination:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **SNMP Destinations** tab in the right panel.
- 4 Locate the SNMP Destination that you want to remove and click the remove icon (✖) on that row.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the SNMP Destination.

Sending Notifications to Syslog Destinations

Syslog Destinations describe how Alert notifications should be sent using the comparatively simple syslog protocol. You need to set up Syslog Destinations before creating Alerts that will use them. Before configuring an Syslog destination, you should be familiar with the information in [“Sending Notifications to Syslog and SNMP Destinations” on page 298](#).

To Add a Syslog Destination:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **Syslog Destinations** tab in the right panel.
- 4 Click the **Add** button.
- 5 Enter parameters:

Parameter	Description
Name	A name for this destination.
Type	UDP or TCP Syslog. This choice cannot be edited later.


- 6 Click **Next**. Enter the secondary parameters:

Parameter	Description
Name	The name for the destination.
Type	This is the value you entered in the previous screen. This value cannot be changed.
Ip/Host	Host name or IP address.
Port	Port (default is 514).
Connection Retry Timeout	(Only for TCP Syslog Destinations) The time, in seconds, to wait before retrying a connection. The default is 5 seconds.


- 7 Click **Save** to create the new Syslog Destination.

To Edit a Syslog Destination:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **Syslog Destinations** tab in the right panel.

- 4 Click the Edit icon (). You can edit the parameters of the Syslog Destination except its type.
- 5 Click **Save** to make the changes, or **Cancel** to return to the Syslog Destination table.

To Remove a Syslog Destination:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Alerts** in the left panel.
- 3 Click the **Syslog Destinations** tab in the right panel.
- 4 Locate the Syslog Destination that you want to remove and click the remove icon () on that row.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Syslog Destination.

Sending Notifications to ESM Destinations

ESM Destinations describe how Alert notifications should be sent to an ArcSight Manager. Set up ESM destinations before creating Alerts that will use them.

If an ArcSight Manager uses a signed SSL certificate, you will need to load it on the Logger.



Note

Audit events for alerts are only written to the Internal Storage Group and not forwarded to ESM by default. If you need to forward the audit events generated for alerts to ESM, please contact customer support for assistance.

To setup Logger to send alerts to an ArcSight Manager:

- 1 If the ArcSight Manager uses a certificate, copy the server SSL certificate file from an ArcSight Console or other component that is already communicating with the target Manager, and upload the certificate file to Logger, as described ["Uploading a Certificate to the Logger:" on page 294](#).



Note

You cannot import the cacerts file, which is a repository of trusted certificates, to the Logger. Instead, you need to import specific SSL certificate files.

- 2 Create an ESM Destination, as described in ["To create an ESM Destination:" on page 292](#).

Scheduled Tasks

Scheduled Tasks displays jobs that are programmed to happen automatically. Job types include Configuration Backup, file transfers, Event Archive, and Saved Searches. The Scheduled Tasks section has three tabs: Scheduled Tasks, Currently Running Tasks, and Finished Tasks.

Make sure you are familiar with the information in ["Impact of Daylight Savings Time Change on Logger Operations" on page 363](#) that can impact a scheduled task.

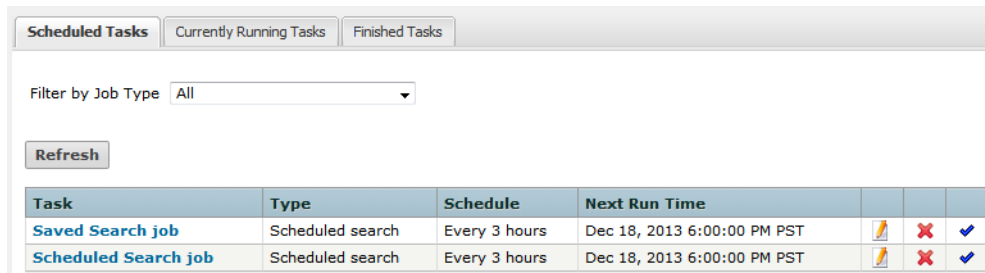
Maximum number of scheduled tasks that can be defined on Logger: No limit.

Scheduled Tasks

The Scheduled Tasks page, shown in [Figure 6-12](#), displays the list of scheduled jobs. Scheduled Tasks can be deleted until the jobs are performed. A drop-down list at the top of the page lets you show All Scheduled Tasks or only tasks of a specific type.

To view Scheduled Tasks:

- 1 Click the **Configuration** (or **Configuration > Settings**) > **Scheduled Tasks**.
- 2 Filter the list by selecting a specific type of Scheduled Task from the drop-down list, or select **All**.
- 3 Click **Refresh** to update the list of tasks.



Task	Type	Schedule	Next Run Time			
Saved Search job	Scheduled search	Every 3 hours	Dec 18, 2013 6:00:00 PM PST			
Scheduled Search job	Scheduled search	Every 3 hours	Dec 18, 2013 6:00:00 PM PST			

Figure 6-12 Scheduled Tasks page

Scheduled Tasks can be created for:

- Saved Search (See [“Scheduled Saved Search”](#) on page 316)
- File Receivers and File Transfer Receivers (See [“Receivers”](#) on page 260)
- Event Archive (See [“Archiving Events”](#) on page 253)
- Configuration Backup (See [“Configuration Backup and Restore”](#) on page 330)

To delete a Scheduled Task:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Scheduled Tasks** in the left panel.
- 3 Click the **Scheduled Tasks** tab in the right panel.
- 4 Locate the Scheduled Task that you want to delete and click the delete icon () on that row.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Scheduled Task.

Currently Running Tasks

The Currently Running Tasks page displays the Scheduled Tasks that are running at the present time. The table shows task name, type, and the date and time that the task started.

To view Currently Running Tasks:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Scheduled Tasks** in the left panel.
- 3 Click the **Currently Running Tasks** tab in the right panel.

- Click **Refresh** to update the list of tasks.

Scheduled Tasks **Currently Running Tasks** Finished Tasks

Filter by Job Type All

Refresh

Task	Type	Start
There are no running tasks to display		

- Filter the list by selecting a specific type of Scheduled Task from the drop-down list, or select **All**.

Finished Tasks

The Finished Tasks page acts like a log of all Scheduled Task runs, with the most recently finished tasks on top.

To view Finished Tasks:

- Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- Click **Scheduled Tasks** in the left panel.
- Click the **Finished Tasks** tab in the right panel.
- Click **Refresh** to update the list of tasks.

Scheduled Tasks Currently Running Tasks **Finished Tasks**

Filter by Result All Filter by Job Type All

Refresh

Task	Type	Start	End	Result	Status
scheduled daily vacuum	Cleaning service	Dec 18, 2013 12:00:00 AM PST	Dec 18, 2013 12:00:10 AM PST	Passed	scheduled vacuum for [2013-12-18] completed
Peer Authorization Expiration Enforcer	Peer authorization expiration	Dec 18, 2013 12:00:00 AM PST	Dec 18, 2013 12:00:00 AM PST	Passed	daily peer authorization expiration enforcement completed
scheduled daily vacuum	Cleaning service	Dec 17, 2013 12:00:00 AM PST	Dec 17, 2013 12:00:00 AM PST	Passed	scheduled vacuum for [2013-12-17] completed
Peer Authorization Expiration Enforcer	Peer authorization expiration	Dec 17, 2013 12:00:00 AM PST	Dec 17, 2013 12:00:00 AM PST	Passed	daily peer authorization expiration enforcement completed

- Filter the list by selecting a specific type of Scheduled Task from the drop-down list, or select **All**.

Filters

The Filters section has two tabs: Filters and Search Group Filters.

Filters Tab

The Filters tab provides a convenient place to manage filters. There are two types of filters:

- **Shared**

A shared filter is visible to all Logger users. Once created, any Logger user can use it to search for events. The query you specify for a shared filter can be a Unified query (that uses keywords, indexed, and non-indexed fields) or a Regex query (that specifies a regular expression). Creating Regex Query shared filters are useful for creating real time alerts, which accept only regex queries.

- **Search Group**

Search group filters provide an access control mechanism to limit the events that users in a particular user group can see. Search Group filters can also be used to limit the events processed by a category of reports (see [“Using Report Category Filters” on page 239](#)). The query for these filters can only contain regular expressions. Only users with administrative privileges can create these filters.

A set of pre-defined filters, also known as system filters, exist on your Logger as well. For more information about system filters, see [“System Filters/Predefined Filters” on page 128](#).

To create a filter:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Filters** in the left panel.
- 3 Click the **Filters** tab in the right panel to create a shared filter, or click the **Search Group Filters** tab to create a search group filter. (See [“Filters Tab” on page 312](#) for information about shared and search group filters.)
- 4 Click **Add** to display the following page.
- 5 Enter a name for the new filter in the Name field.
Filter names are case-sensitive.
- 6 If you are creating a shared filter, select **Unified** or **Regex Query**.
For Search Group filters, select **Search Group**.



Note

Non-administrator users cannot create Search Group filters.

- 7 Click **Next**.
- 8 If you selected Unified or Regex Query method in the previous step, enter the query for the new filter.

For Unified queries:

When you type a query, Logger's Search Helper enables you to quickly build a query expression by automatically providing suggestions, possible matches, and applicable operators. See [“Search Helper” on page 103](#) for more information.

OR

Click **Advanced Search** to use the Search Builder Tool to create the query. For details about using the Search Builder Tool, see ["Using the Advanced Search Builder Tool" on page 96](#).

For Regex queries:

Enter the regular expression in the Query text box.

- 9 Click **Save**.



If you created a Search Group filter, make sure that you associate it to a user group, as described in ["Search Group Filters Tab" on page 314](#).

To create a filter by copying an existing filter:

- 1 Click the **Configuration** tab, click **Filters** in the sub-menu, and then click the **Filters** tab.
- 2 Locate the filter to copy from the list of filters on the Filters tab. Click the copy icon ().

A new filter with the name "Copy of <filtername>" is created.

- 3 Change the name of the filter and edit the query for the new filter if necessary.
- 4 Click **Save**.

To edit a filter:

- 1 Click the **Configuration** tab, click **Filters** in the sub-menu, and then click the **Filters** tab.
- 2 Find the filter that you want to edit and click the Edit icon () on that row.
- 3 Change the information in the form and click **Save**.

To delete a filter:

- 1 Click the **Configuration** tab, then click **Filters** in the sub-menu.
- 2 Find the filter that you want to delete in the table.
- 3 Click the Delete icon (). Confirm the delete.

Search Group Filters Tab

The Search Group Filters tab is used to manage the association of User Groups with Search Group Filters.

Filters		Search Group Filters
<p>You may assign a search filter to a search group that will be appended to all searches performed by users in that search group.</p> <p>To create a new search group filter, you must first go to the Filters page and add a new filter of type Search Group.</p>		
Name	Filter	Description
Default Logger Search Group	NONE	The default search group allows both local and distributed searches.

Figure 6-13 Search Group Filters Tab



In the Search Group Filters tab (shown in [Figure 6-13](#)), the User Group of type Search Group is listed in the left column and the associated filter is listed in the middle column.

Search Group Filters can be used to restrict events in the following two ways:

- **Restrict the events processed by a Report Category**—A Search Group Filter can be associated directly with a Report Category. This association provides a way to restrict the events processed by all the reports in a Report Category.

When a Search Group filter is used to restrict the events processed by a Report Category, you do not need configure the Search Group in the Search Group Filters tab as described below. After creating the filter (of type Search Group), you can go directly to the Reports Category Filters page of the Report tab and select the filter for the Report Category. For more information, see [“Using Report Category Filters” on page 239](#).

- **Restrict the events visible by members of a user group**—A Search Group Filter can be associated with a user group (of type Logger Search). This association means that all members of the user group only see events which match the Search Group Filter. User groups (described in more detail later in this chapter) provide a way of assigning privileges to a specified set of users.



Users who belong to a User Group that does not have a Search Group Filter will see all events.


To create, edit, or delete Search Group Filters, see [“Filters” on page 312](#). To create, edit, or delete User Groups, see [“Users/Groups” on page 392](#).



Only users that are members of a System Admin group can assign Search Group Filters. For more information, see [“Users/Groups” on page 392](#).

To associate a Search Group Filter with a User Group:

- 1 If the User Group that you want to associate with the Search Group Filter does not exist, create a new User Group of type Search Group. For instructions, see [“Users/Groups” on page 392](#).

- 2 If the Search Group Filter you want to associate with the User Group does not exist, create a filter of type Search Group. For instructions, see [“To create a filter:” on page 312](#). When creating the filter, from the **Type** pull-down menu select the **Search Group** option.
- 3 Click the **Configuration** tab, click **Filters** in the sub-menu, and then click the **Search Group Filters** tab. The page shown in [Figure 6-13](#) is displayed.
- 4 Find the User Group to which to apply a Search Group Filter. Click the edit icon ().
- 5 Select a filter from the pulldown list. (Only Search Group type filters are listed.)
- 6 Click **Save**.

Saved Searches

A Saved Search, like a Saved Filter, recalls a specific query. A Saved Search includes a time range, unlike a Saved Filter, which supports the creation of scheduled event reporting. Also, a saved filter does not include the field set information that determines the fields that are displayed for each event in the search results. For information about Saved Search Alerts, see [“Alerts” on page 296](#).

You can schedule a saved search to run at a specific interval. For more information, see [“Scheduled Saved Search” on page 316](#).

Make sure you are familiar with the information in [“Impact of Daylight Savings Time Change on Logger Operations” on page 363](#) before adding a Saved Search.

Saved Searches

The Saved Searches tab displays all Saved Searches and supports Adding, Editing, and Deleting Saved Searches.

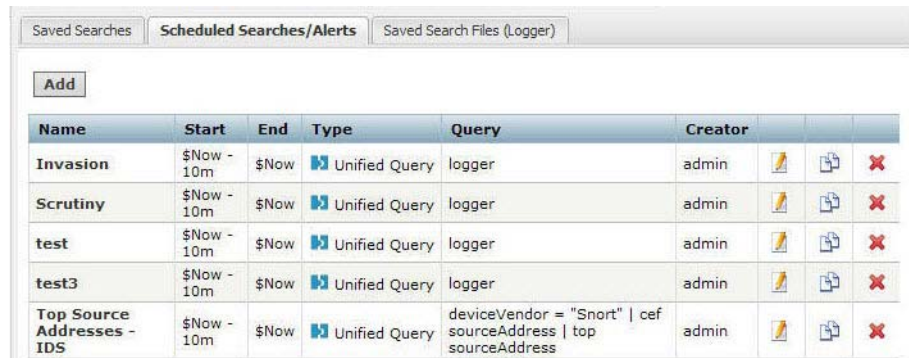
To add a Saved Search:

- 1 Click the **Configuration** (or **Configuration > Settings**) > **Saved Search**.
- 2 Click **Add** and enter the following parameters:

Parameter	Description
Name	A name for this Saved Search. This name will be used for exported output files, with the Saved Search date and time appended.
Start Time	Absolute date and time of earliest possible event. Alternatively, check Dynamic to specify the start time relative to the time when the Saved Search job is run.
End Time	Absolute or Dynamic date and time of latest possible event, as described above.
Query Terms	Enter the query in the text field or select one or more Filters from the list below the text field. When you type a query, Logger’s Search Helper enables you to quickly build a query expression by automatically providing suggestions, possible matches, and applicable operators. See “Search Helper” on page 103 for more information.

Parameter	Description
Local Search	Check this box to limit the Saved Search to the local Logger box. If the Local Search box is left unchecked, the Saved Search will include all Peer Loggers as well as the local Logger.

- 3 Click **Save** to add the new Saved Search, or **Cancel** to quit.



The screenshot shows a web interface with three tabs: "Saved Searches", "Scheduled Searches/Alerts", and "Saved Search Files (Logger)". The "Scheduled Searches/Alerts" tab is active. Below the tabs is an "Add" button. A table lists several saved searches with columns for Name, Start, End, Type, Query, and Creator. Each row also has icons for edit, delete, and a status icon.

Name	Start	End	Type	Query	Creator			
Invasion	\$Now - 10m	\$Now	Unified Query	logger	admin			
Scrutiny	\$Now - 10m	\$Now	Unified Query	logger	admin			
test	\$Now - 10m	\$Now	Unified Query	logger	admin			
test3	\$Now - 10m	\$Now	Unified Query	logger	admin			
Top Source Addresses - IDS	\$Now - 10m	\$Now	Unified Query	deviceVendor = "Snort" cef sourceAddress top sourceAddress	admin			

Figure 6-14 Saved Search page

To edit a Saved Search:

- 1 Click the **Configuration** (or **Configuration > Settings**) > **Saved Search**.
- 2 Find the Saved Search that you want to edit and click the Edit icon () on that row.
- 3 Change the information in the form and click **Save**.

To delete a Saved Search:

- 1 Click the **Configuration** (or **Configuration > Settings**) > **Saved Search**.
- 2 Find the Saved Search that you want to delete in the table.
- 3 Click the Delete icon (). Confirm the delete.

Scheduled Saved Search

A scheduled Saved Search schedules a Saved Search to be run at a later time. Before you schedule a Saved Search, you must have created or saved at least one Saved Search. A scheduled Saved Search can be also configured to generate an alert. For more information about scheduled Saved Search Alerts, see [“Creating a Saved Search Alert” on page 303](#).

The results of a scheduled Saved Search are written to a file, as described in [“Saved Search Files” on page 320](#).

To add a scheduled Saved Search:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.
- 4 Click **Add**. The screen shown in [Figure 6-15](#) is displayed.

Name

Schedule

Everyday

Hour of day

Hours (24 hour format)

Saved Searches

Invasion
Scrutiny
test
test3
Top Source Addresses - IDS

Use ctrl-click to select or deselect items

Job type

Search

Search Result Export Options

Export Options

☒ Export to remote location

☐ Save to Logger

File format

PDF

Export directory name

Title

Fields

Event Time, Receipt Time, Device, Logger, Name, Version, Device Vendor, Device Product, Device Version, Signature ID, Severity

☒ All fields

Chart type

Column

Chart result limit

10

Include summary

☐

Include only CEF events

☐

Save

Cancel

Figure 6-15 Saved Search Jobs page

5 Enter the following parameters:


Parameter	Description
Name	A name for this Scheduled Saved Search Job.
Schedule	<p>Choose Everyday or Days of Week from the first pulldown menu.</p> <p>If Everyday, select Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the text box.</p> <p>If Days of Week, enter the days (day 1 is Sunday) in the text box. Then choose Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the second text box.</p> <p>For example, to perform the Saved Search every day at 2 a.m., select Everyday in the first pulldown menu, then choose Hour of Day from the second pulldown menu and enter 2 in the text box. To perform the Saved Search every day at 2 a.m. and 3 p.m., enter 2,15 in the text box.</p> <p>For another example, to perform the Saved Search Tuesdays and Thursdays at 10 p.m., select Days of Week from the first pulldown menu and enter 3,5 for days. Then choose Hour of Day from the second pulldown menu and enter 22 in the text box.</p>

Parameter	Description
Saved Searches	Select from the list of Saved Searches. If a Saved Search to suit your needs does not exist in the list, click the Saved Searches tab (to the left of Scheduled Searches tab) to define it. For more information about defining a Saved Search query, see “Saved Searches” on page 315 . When <i>multiple</i> Saved Searches are specified in one scheduled saved search job, the resulting file contains the number of hits for each Saved Search and not the actual events.
Job Type	Select Search for a scheduled Saved Search.
Export Options	<p>For the Logger appliance: Select from one of these options:</p> <ul style="list-style-type: none"> Export to remote location—The file is written to an NFS mount, a CIFS mount, or a SAN system. Save to Logger—The file is saved to the Logger’s onboard disk. If the file is saved locally, use the Saved Search Files (“Saved Search Files” on page 320) feature to access those files. <p>For the software version of Logger: The only applicable option is “Save to Logger”, which is preselected for you.</p>
File Format	<p>CSV, for comma-separated values file.</p> <p>PDF, for a report-style file that contains search results as charts and in tables. You must specify a title for the report in the Title field. If the search query contains an operator that creates charts such as chart, top, and so on, charts are included in the PDF file. In that case, you can also set the Chart Type and Chart Result Limit fields. These fields are described later in this table.</p>
Export Directory Name	<p>For the Logger appliance, select the directory where the search results will be exported from the pulldown menu.</p> <p>For the software version of Logger, enter the directory path in this field, which can be a path to a local directory or to a mount point on the machine on which the software version of Logger is installed.</p> <p>If a directory of the specified name does not exist, it is created. If a directory of the specified name exists and the Overwrite box is not checked, an error is generated. If the Overwrite box is checked, the existing directory contents are overwritten.</p>
Title	(Optional) A meaningful name that appears on top of the PDF file. If no title is specified, “Untitled” is included.
Fields	<p>A list of event fields that will be included in the exported file.</p> <p>By default, all fields are included.</p> <p>You can enter fields or edit the displayed fields by deselecting All Fields.</p>
Chart Type (for PDF only)	<p>(If the search query includes an operator that creates a chart, this field is meaningful; otherwise, it is ignored.)</p> <p>Type of chart to include in the PDF file. You can select from: Column, Bar, Pie, Area, Line, Stacked Column, Stacked Bar.</p> <p>Note: If the Chart Type is different from the chart displayed on the Search Results screen, the value selected for this option overrides the one shown in the screen. Therefore, the exported PDF contains the chart you specify for this option and not the one shown on the screen.</p>


Parameter	Description
Chart Result Limit (for PDF only)	<p>(If the search query includes an operator that creates a chart, this field is meaningful; otherwise, it is ignored.)</p> <p>Number of unique values to plot. Default: 10</p> <p>If the configured Chart Result Limit is less than the number of unique values for a query, the top values equal to the Chart Result Limit are plotted. That is, if the Chart Result Limit is 5 and 7 unique values are found, the top 5 values will be plotted.</p>
Include Summary	Check this box to include an event count with the Saved Search, or a total when more than one Saved Search is specified.
Include Non-CEF Events	<p>Check this box to include all events. Uncheck the box to include only CEF events in the output.</p> <p>For more information about the Common Event Format (CEF), refer to Implementing ArcSight CEF. For a downloadable a copy of this guide, search for "ArcSight Common Event Format (CEF) Guide" on the Protect 724 Community at https://protect724.arcsight.com.</p>

- 6 Click **Save** to add the new scheduled Saved Search, or **Cancel** to quit.

To edit a scheduled Saved Search:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.
- 4 Locate the Saved Search Job that you want to edit and click the Edit icon () on that row.
- 5 Change the parameters of the Saved Search Job.
- 6 Click **Save** to update the Saved Search Job, or **Cancel** to abandon your changes.

To delete a scheduled Saved Search:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Saved Search** in the left panel.
- 3 Click **Scheduled Searches** in the right panel.
- 4 Locate the Saved Search Job that you want to delete and click the delete icon () on that row.
- 5 Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Saved Search Job.

Saved Search Files

Access Saved Search results that were saved to Logger with the Saved Search Files command. Saved Search Files can be retrieved (streamed to the browser) or deleted. Click Refresh to update the list of files.



<div> <div>Saved Searches</div> <div>Scheduled Searches/Alerts</div> <div>Saved Search Files</div> </div>						
<div>Refresh</div>						
Name	Last Modified	Size	State	Error Message		
Exported_Search_Results.csv	Dec 18, 2013 3:16:23 PM PST	3.4 KB	Exported			

Figure 6-16 Saved Search Files page

Search

The search screen enables you to:

- Add search indexes for field query search operations
- Tune advanced search options
- View and delete custom field sets
- View default schema
- View custom schema fields
- View and end currently running search tasks
- View and add parsers for specific log types

For general search information, including how to search, see [“Searching and Analyzing Events” on page 75](#).

Adding Search Indexes

See [“Indexing” on page 121](#) for more information.

Tuning Advanced Search Options

The advanced search options support internationalization (i18n) choices. To change these options, click **Configuration** (or **Configuration > Settings > Search > Search Options**).

Search Indexes	Search Options	Fieldsets	Default Fields	Custom Fields	Running Tasks	Parsers
Edit Search Options						
Most users shouldn't need to adjust these settings						
Field Search Options						
Case sensitive		<input type="text" value="Yes"/>				
Full-text Search Options						
Use primary delimiters		<input type="text" value="Yes"/>				
Use secondary delimiters		<input type="text" value="No"/>				
Regular Expression Search Options						
Case sensitive		<input type="text" value="No"/>				
Unicode case sensitive		<input type="text" value="No"/>				
Check for canonical equality		<input type="text" value="No"/>				
Search Display Options						
Populate rawEvent field for syslog events		<input type="text" value="No"/>				
Show source and sourceType fields		<input type="text" value="No"/>				
Field Summary Options						
Use Field Summary		<input type="text" value="Yes"/>				
Discover fields		<input type="text" value="No"/>				

The following table lists the advanced search options you can view and configure. Several of the options on this screen will require you to reboot your Logger appliance or restart your software Logger.

Option	Description
Field Search Option	
Case sensitive	<p>Default: Yes</p> <p>Controls whether to differentiate between upper- and lower-case characters during a search. When this option is set to No, searching for "login" will find "login," "Login," and "LOGIN".</p> <p>You must reboot the Logger appliance/restart the software Logger for this change to take effect.</p> <p>Notes:</p> <ol style="list-style-type: none"> 1 Case-sensitive search only applies to the local Logger. Peer Loggers will continue to use case-insensitive search. 2 Full-text search (keyword search) is case insensitive. You cannot change its case sensitivity. 3 Set this option to Yes to increase local query performance.
Full-text Search Options	
Use primary delimiters	<p>Default: Yes</p> <p>Controls whether primary delimiters are applied to an event to tokenize it for indexing.</p> <p>A primary delimiter tokenizes an event for indexing. For example, an event "john doe the first" is tokenized into "john" "doe" "the" "first" using the "space" primary delimiter.</p> <p>Supported primary delimiters: space, tab, newline, comma, semi-colon, (,), [,], {, }, ", , *, >, <, !</p>
Use secondary delimiters	<p>Default: No</p> <p>Controls whether secondary delimiters are applied to an event to further tokenize a token created by a primary delimiter thus enabling searches that can match a part of a primary token.</p> <p>For example, you can search for "hp.com" in http://www.hp.com.</p> <p>Supported secondary delimiters: =, ., :, /, \, @, -, ?, #, \$, &, _, %</p>
Regular Expression Search Options	
Case sensitive	<p>Default: No</p> <p>See "Case sensitive" on page 322.</p> <p>You must reboot the Logger appliance/restart the software Logger for this change to take effect.</p>
Unicode case sensitive	<p>Default: No</p> <p>Controls whether events in languages other than English should be compared in a case-sensitive way.</p> <p>Caution: HP strongly recommends that you do not change this option.</p> <p>You must reboot the Logger appliance/restart the software Logger for this change to take effect.</p>

Option	Description
Check for canonical equality	<p>Default: No</p> <p>Controls whether events in languages other than English should be compared using locale-specific algorithms.</p> <p>Caution: HP strongly recommends that you do not change this option.</p> <p>You must reboot the Logger appliance/restart the software Logger for this change to take effect.</p>
Search Display Options	
Populate rawEvent field for syslog events	<p>Default: No</p> <p>Controls whether raw events are displayed in a formatted column called rawEvent using the Raw Event field set. This option applies to syslog events only. If you want to view the raw events associated with CEF events, you do not need to configure this setting. Instead, configure the connector that is sending events to Logger to populate the rawEvent field with the raw event.</p> <p>Note: Even though the rawEvent column displays the raw event, this column is not added to the Logger database and is not indexed. Therefore, you can only run a keyword (full-text) or regular expression search on the event.</p>
Show Source and SourceType fields	<p>Default: No</p> <p>Controls whether the Source and SourceType fields are included in the Field Summary and query results.</p> <p>You must reboot the Logger appliance/restart the software Logger for this change to take effect.</p> <p>Note: Setting this option to Yes can impact query performance.</p>
Field Summary Options	
Use Field Summary	<p>Default: Yes</p> <p>Controls the whether the Field Summary panel is included in the search results by default. Regardless of the default, you can change the setting on-the-fly by using the Fields Summary checkbox on the Search screen.</p>
Discover Fields	<p>Default: No</p> <p>Controls whether the Field Summary feature automatically detects non-CEF fields in raw events. Regardless of the default, you can change the setting on-the-fly by using the Discover Fields checkbox on the Search screen.</p> <p>This field is hidden if Use Field Summary is set to No.</p>

Viewing and Deleting Field Sets

You can view the field sets you have created and the predefined field sets on the Fieldsets tab (**Configuration** (or **Configuration > Settings**) > **Search > Fieldsets**). You can also delete the field sets you created.



Note

- You need to have the "Edit, save, and remove fieldsets" privilege to delete a custom field set.
- You can only delete the field sets you create, and not the predefined ones available on Logger.

To delete a custom field set:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Search** from the left panel and then click Fieldsets.

- 3 In the Fieldsets tab, identify the field set you want to delete and click the delete (✖) icon.
- 4 Confirm the deletion.

Viewing Default Fields

The Logger schema comes with a set of predefined fields. Some of these fields are already indexed for improved search speed and efficiency. You can add custom fields to the Logger schema and index them for field-based search. A field-based search can only use fields in Logger's schema.



The size of each field in the Logger schema is predetermined. If the string you are searching for is longer than the field-length, you should use a STARTSWITH rather than an = search, and include no more than the number of characters in the field size. For more information, see ["Field-based Search" on page 80](#).

The Logger Default Fields tab (**Configuration** (or **Configuration > Settings**) > **Search > Default Fields**) displays the predefined fields included in the Logger schema. It includes the Display Name, Type, Length, and Field Name for each default field. To view information on existing custom fields, see ["Viewing Custom Fields" on page 324](#).

To view the default schema fields:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Search** from the left panel and then click **Default Fields**.

Search Indexes	Search Options	Fieldsets	Default Fields	Custom Fields	Running Tasks	Parsers
Display Name	Type	Length	Field Name	Indexed		
agentAddress	TEXT	16	agt			
agentHostName	TEXT	40	ahost			
agentNtDomain	TEXT	40	agentNtDomain			
agentSeverity	TEXT	-	agentSeverity			✓
agentType	TEXT	16	at			✓
agentZone	TEXT	200	agentZone			
agentZoneName	TEXT	50	agentZoneName			
agentZoneResource	TEXT	100	agentZoneResource			
agentZoneURI	TEXT	2048	agentZoneURI			
applicationProtocol	TEXT	40	app			✓
baseEventCount	LONG	-	cnt			✓
bytesIn	LONG	-	in			✓
bytesOut	LONG	-	out			✓
categoryBehavior	TEXT	200	categoryBehavior			✓
categoryDeviceGroup	TEXT	200	categoryDeviceGroup			✓
categoryObject	TEXT	100	categoryObject			✓
categoryOutcome	TEXT	100	categoryOutcome			✓
categorySignificance	TEXT	100	categorySignificance			✓
categoryTechnique	TEXT	200	categoryTechnique			✓

- 3 The Default Fields tab displays the default fields. You can sort the fields by clicking the column headers.

Viewing Custom Fields

You can view the custom fields that have been added to the Logger schema under the Custom Fields tab (**Configuration** (or **Configuration > Settings**) > **Search > Custom Fields**).

Search Indexes	Search Options	Fieldsets	Default Fields	Custom Fields	Running Tasks	Parsers
----------------	----------------	-----------	----------------	----------------------	---------------	---------

Display Name	Type	Length	Field Name	Actual Field Name	Creator	Created
testbigint	BIGINT	-	testbigint	ad.testbigint.i	admin	Feb 28, 2012 8:45:23 AM PST
testdatetime	DATETIME	-	testdatetime	ad.testdatetime.d	admin	Feb 28, 2012 8:45:37 AM PST
testdouble	DOUBLE	-	testdouble	ad.testdouble.r	admin	Feb 28, 2012 8:45:10 AM PST
testtext	TEXT	255	testtext	ad.testtext	admin	Feb 28, 2012 8:45:58 AM PST

This page lists all custom schema fields that have been saved. You can view the alphabetical list of fields, but cannot edit or delete them.

For detailed information about custom schema fields, see [“Adding or Importing Schema Fields” on page 345](#).

Running Search Tasks

The Running Tasks page displays the search tasks that are running at the present time. The table shows the session ID, user who started the tasks, the date and time that the task started, the number of hits, the number of scanned events, the elapsed time, and the query.

Search Indexes	Search Options	Fieldsets	Default Fields	Custom Fields	Running Tasks	Parsers
----------------	----------------	-----------	----------------	---------------	----------------------	---------

Refresh						
Session ID	User	Start	Hits	Scanned	Elapsed	Query
104857621	admin	Dec 18, 2012 2:40:29 PM PST	28,408	31,501	00:04.224	_deviceGroup in ["Logger Internal Event Device"]
104857623	guest	Dec 18, 2012 2:40:31 PM PST	11,245	13,901	00:02.249	_deviceGroup in ["Logger Internal Event Device"]

Once a task finishes, the task's entry on the Running Tasks page is removed. (The task entry is removed upon page refresh, either when you refresh the browser page or when you navigate away from this page and come back to it.)


To view Running Tasks:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Search** in the left panel.
- 3 Click the **Running Tasks** tab in the right panel.
The Running tasks are displayed.

Ending Currently Running Tasks

You might need to end a currently running search task when it is taking too long to run, or appears to be stuck and slowing the overall Logger performance. When a search initiated as a result of any of the following operations is **in-progress**, the Running Tasks page (**Configuration** (or **Configuration > Settings**) > **Search > Running Tasks**) displays the currently running process.

- A manual search on local or peer Logger (**Analyze > Search**)
- A scheduled search (**Configuration** (or **Configuration > Settings**) > **Saved Search > Scheduled Searches/Alerts**)
- A saved search alert (**Configuration** (or **Configuration > Settings**) > **Saved Search > Scheduled Searches/Alerts**)
- A search export, with the “Rerun query” option checked (**Analyze > Search > Export Results**)

To end a process, click the  icon. You must have admin user privileges to end a running search process.

View and Add Parsers for Specific Log Types

The Parsers tab on this screen is the same as the one documented in the Event Input section. See [“Parsers” on page 278](#) for information on parsers and how to add them.

Peer Loggers

Logger can establish peer relationships with one or more Loggers or ArcSight Managers to enable distributed searches. To search other Loggers or Managers, you must define one or more peers.

When two systems peer with each other, one initiates the relationship. The initiator sends credentials to authenticate itself to the target system. If the authentication succeeds, a peer relationship is established between the two systems.

Adding a peer on a Logger is a bi-directional process. That is, when Logger A adds peer access for Logger B, Logger B automatically adds peer access for Logger A. Similarly, if you delete the peer access for B on A, the peer access for A is automatically deleted on B.



Note

On a Logger using local or RADIUS authentication, you can use either authentication method, although peer authorization ID and code are recommended for reasons described below. However, if you are using SSL Client Authentication (CAC) on your Logger, authorization ID and code is the only way to authenticate a peer. You cannot use a user name and password.

FIPS-enabled Loggers are not limited to a specific authentication method. Therefore, you can use any listed below.

Peer Loggers can authenticate using any of these methods:

- User name and password
A user name configured on the Logger is used for authentication
- Peer Authorization ID and Code
Authorization ID and Code generated on a remote Logger are used by the initiator Logger to peer with it. The generated ID and Code are specific to the initiator Logger because the IP address of the initiator is used to generate the ID and code, and can be used only for peering from the initiator. Therefore, this method is more secure than using user name and password.

Guidelines

You should be aware of these guidelines when peering Loggers.

- Logger 5.5 can peer with ESM 6.5c, Logger 5.3 SP1, and Logger 5.3.
- You can configure a maximum of 20 peers for a Logger.
- The time and date on the system on with the software Logger is installed must be set correctly with respect to its timezone to peer with other Loggers. HP recommends that you configure the Logger system to synchronize its time with an NTP server regularly.
- If the remote Logger is configured for SSL Client authentication (CAC), you must configure an authorization ID and code on the initiator Logger.

There are no special authentication requirements for FIPS-enabled Loggers. Such Loggers can use any of the allowed authentication methods.

- Peer Loggers cannot be edited, however you can delete and re-add a peer.
- If you are running distributed searches (searches across peers), follow these additional guidelines:
 - ◆ A user must belong to the Logger Search User Group with “Search for events on remote peers” privilege set to Yes and the Logger Rights Group with “View registered peers” privilege set to Yes. See [“Searching Peer Loggers \(Distributed Search\)” on page 108](#).
 - ◆ Users performing search operations on peers have the same privileges on the peer that they have on the Logger they are logged in.
 For example, User A is restricted by a search group filter to only search for events in which deviceVendor is set to “Cisco”. When User A performs a search operation across Logger A’s peers, the same constraint (to search events where deviceVendor = “Cisco”) is applied on all peers.
- If you are running distributed reports (reports across peers), see [“Selecting Device Groups, Storage Groups, Devices, or Peers” on page 177](#).
- If user name and password are used for authenticating to a remote peer Logger, the credentials are only used one-time, during the peering relationship set up. After a relationship has been established, the credentials are not saved (on the Peer Loggers page) and the peers do not authenticate periodically. Therefore, if the user name or password used to establish a relationship is changed at a later date or the user name is deleted, peering relationship is not broken.
- If you delete the peering relationship, you will need to re-add the peer re-establish the relationship.
- If a peering relationship breaks, for example, if you need to reinstall a peer Logger, you will need to delete and re-add the peer re-establish the relationship.
- The following example illustrates the steps you need to follow to set up peering between two Loggers.

Logger A

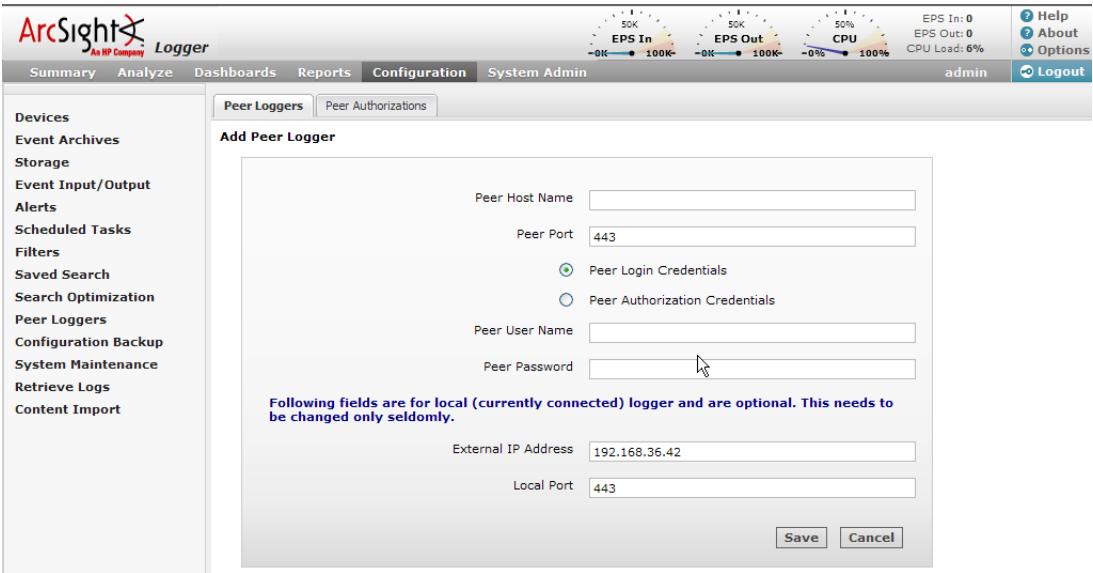
Logger B

- 1 Select the Logger that will initiate the establishment of the peering relationship.
 In this example, Logger A will initiate the relationship.
 - 2 If Logger B is configured to use user name and password authentication, go to [Step 3](#).
 If Logger B is configured to use SSL Client Authentication (CAC), go to [Step 4](#).
 - 3 Set up a user name and password that Logger A will use to authenticate itself when peering with this Logger, as described in [“Users/Groups” on page 392](#).
 - 4 Generate an Authorization ID and Code that Logger A will use for authenticating to Logger B, as described in [“To generate Authorization ID and Code for configuring a peer relationship:” on page 330](#).
-

Logger A

Logger B

- 5 Add Logger B's information, as described in ["To add a peer Logger:" on page 328](#):
- If Logger B uses user name and password, use the user name and password you configured in [Step 3](#).
- If Logger B uses SSL Client Authentication, use the Authorization ID and Code you generated in [Step 4](#).



To add a peer Logger:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Peer Loggers** from the left panel.
- 3 Click **Add** and enter the following parameters.

Parameter	Description
Peer Host Name	The remote Logger's hostname or IP address.
Peer Port	443, by default for Logger appliance and pre-5.0 software Loggers. Port you configured when installing software Logger. See "Guidelines" on page 326 .

Parameter	Description
Peer Login Credentials	Select Peer Login Credentials for password-based authentication with the remote Logger.
Peer Authorization Credentials	<p>Select Peer Authorization Credentials for SSL client authentication with the remote Logger. (See “SSL Client Authentication” on page 386.)</p> <p>If the peer-relationship initiating Logger has version 3.0.x installed and the other Logger is running 4.0 GA with SSL Client Authentication enabled, enter the generated Authorization ID in the User Name field and the Code in the Password field on the 3.0.x Logger.</p>
If you selected Peer Login Credentials...	
Peer User Name	The user name to use when connecting to the remote Logger.
Peer Password	The password for the user on the remote Logger.
If you selected Peer Authorization Credentials...	
Peer Authorization ID	Enter the authorization ID of the other Logger to which this Logger is initiating a peering relationship. (See “To generate Authorization ID and Code for configuring a peer relationship:” on page 330 for more information.)
Peer Authorization Code	Enter the authorization code of the other Logger to which this Logger is initiating a peering relationship. (See “To generate Authorization ID and Code for configuring a peer relationship:” on page 330 for more information.)
These fields need to be updated in rare circumstances. For more information, read the description of each field in this table.	
External IP Address	<p>In most cases, the value in this field matches the IP address you use to connect to this Logger from your browser, and you do not need to do anything.</p> <p>However, if the IP address does not match (for example, when the Logger is behind a VPN concentrator), change the value to match the IP address with which you connect to this Logger.</p>
Local Port	Make sure the value of this field is set to 443.

- 4** Click **Save** to add the new Logger, or **Cancel** to quit.

To delete a peer Logger:

- 1** Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2** Click **Peer Loggers** from the left panel.
- 3** Locate the Peer that you want to delete and click the delete icon (✖) on that row.
- 4** Confirm the deletion by clicking **OK**, or click **Cancel** to retain the Peer.

To view peers of a Logger:

A list of remote Loggers that a Logger peers with is displayed on the Peer Loggers page (**Configuration** (or **Configuration > Settings**) > **Peer Loggers**).

Authorizing Peers

Use the following procedure to generate the Authorization ID and Code on the target Manager or Logger with which you want to establish a peer relationship. (Logger B in the example described earlier in this section.) After that, use the ID and Code on the initiating Manager or Logger when configuring the peer relationship. (Manager Logger A in that example.)

To generate Authorization ID and Code for configuring a peer relationship:

- 1 Click **Configuration** (or **Configuration > Settings**) from the top-level menu bar.
- 2 Click **Peer Loggers** from the left panel.
- 3 In the **Peer Authorizations** tab, click **Add**.
- 4 Enter the hostname for the peer Logger and the port (if using a non-default port).
- 5 Click **Save**.

The authorization ID and authorization Code are displayed. Copy and paste this information when adding a peer Logger that is configured to use SSL client authentication.

Configuration Backup and Restore

By default, Logger does not back up any content. However, you can configure it to backup the following content to a remote system:

- All non-event data
- Reports content only

You can back up this content on ad-hoc basis or schedule it to run periodically. The content is saved to the backup location in a single .tar.gz format file.



Ensure that both Configuration Backups (for configuration settings) and Event Archives (for data) run on a regular basis and are stored in a remote location. In the event of catastrophic failure, you will need to restore the most recent Configuration Backup and Event Archive. For information on Event Archives, see [“Event Archives” on page 250](#).

The following table lists the information included in the backup when you back up all non-event data and reports-only data.

All non-event data backup includes...	Reports-only backup includes...
System Information Logs Global Settings User and Group Information All Configuration Settings Existing Filters and Saved Searches Logger Monitor settings The following Reports content: <ul style="list-style-type: none"> Queries, Reports, Parameters, Parameter Value Groups, Dashboards Templates 	The following Report content only: <ul style="list-style-type: none"> Queries, Reports, Parameters, Parameter Value Groups, Dashboards Templates

You can use the backed up content to:

- Restore a Logger that is not functioning as expected or that has been reset to its factory defaults
- Copy content from one Logger to another



When you restore content to a Logger, the existing content on it is deleted.

Caution


Running a Configuration Backup (Ad-hoc or Scheduled)

The screenshot displays the ArcSight Logger web interface. At the top, there are status gauges for EPS In, EPS Out, and CPU Load. Below these is a navigation bar with tabs: Summary, Analyze, Dashboards, Reports, Configuration, System Admin, and a user profile section (admin, Logout). The left sidebar lists various system components. The 'Configuration' tab is active, showing the 'Configuration Backup' section. Within this section, the 'Edit Configuration Backup' form is visible. The form includes the following fields and options:

- Protocol:** A dropdown menu set to 'SCP'.
- Port:** A text input field containing '22'.
- Ip/Host:** An empty text input field.
- User:** An empty text input field.
- Password:** An empty text input field.
- Remote directory:** An empty text input field.
- Backup content:** A dropdown menu set to 'All'.
- Schedule:** A checkbox labeled 'One time only' which is checked.

At the bottom right of the form are 'Save' and 'Cancel' buttons.

To run a configuration backup or to edit the configuration backup settings:

- 1 Click the **Configuration** (or **Configuration > Settings**) > **Configuration Backup**.
- 2 Click the () icon and enter the following parameters

Parameter	Description
Protocol	SCP
Port	The port on which the Logger should connect to the remote system
Ip/Host	The IP address or hostname of the remote system
User	A user on the remote system with write privileges on the backup folder (specified in Remote Directory, below)
Password	Password for the user Note: The password cannot contain these characters: % = ; " ' < >
Remote Directory	The folder on the remote system in which to save the configuration backup files
Backup Content	Whether to backup all non-event data or only the report content Select All for all non-event data or Report Content Only for only the report content.
Schedule	<p>If you check One Time Only, other fields are hidden and the Configuration Backup occurs just once (ad-hoc), when you click Save.</p> <p>Choose Everyday or Days of Week from the first pulldown menu.</p> <p>If Everyday, select Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the text box.</p> <p>If Days of Week, enter the days (day 1 is Sunday) in the text box. Then choose Hour of Day or Every from the second pulldown menu. Enter the hours (1-23) in the second text box.</p> <p>For example, to backup every day at 2 a.m., select Everyday in the first pulldown menu, then choose Hour of Day from the second pulldown menu and enter 2 in the text box. To backup every day at 2 a.m. and 3 p.m., enter 2,15 in the text box.</p> <p>For another example, to backup Tuesdays and Thursdays at 10 p.m., select Days of Week from the first pulldown menu and enter 3,5 for days. Then choose Hour of Day from the second pulldown menu and enter 22 in the text box.</p> <p>Make sure you are familiar with the information in "Impact of Daylight Savings Time Change on Logger Operations" on page 363.</p>

- 3 Click **Save**.

If you chose to run the backup One Time Only, it is run right away. Otherwise, it is scheduled to run at the specified time.

Restoring from a Configuration Backup

Make sure you are familiar with these guidelines before you restore a backup file on Logger:

- When you restore content to a Logger, the existing content on it is deleted.

Logger restores the specific environment settings that were current at the time the backup was taken. Any configuration settings that were updated between the time of the backup and the time of the restore are lost.

- You must restore the content to the same version of Logger that was used to create the backup file.
- You must restore to the same form of Logger (software, appliance, or VMWare.)
- For appliance Loggers, the Logger appliance model must be the same as the one used to create the backup file.
- For software Loggers, the operating system that Logger is running must be the same as the one used to create the backup file.

To restore from a configuration backup:

- 1 Click the **Configuration** (or **Configuration > Settings**) > **Configuration Backup**.
- 2 Click **Restore**.
- 3 Click **Browse** to locate the backup file.
- 4 Click **Submit** to start the restore process.

Editing Configuration Backup Settings


See [“Running a Configuration Backup \(Ad-hoc or Scheduled\)”](#) on page 331.

System Maintenance

Certain operations on Logger, such as database defragmentation, extending the storage volume size, adding storage groups, and adding additional schema fields, require that Logger be in a maintenance state—a state in which operations related to data on the Logger are not running. Maintenance mode enables you to place the Logger in such a state. When a Logger is in maintenance mode:

- Events are not processed
- Reports are not generated
- Search cannot run
- Scheduled jobs do not run

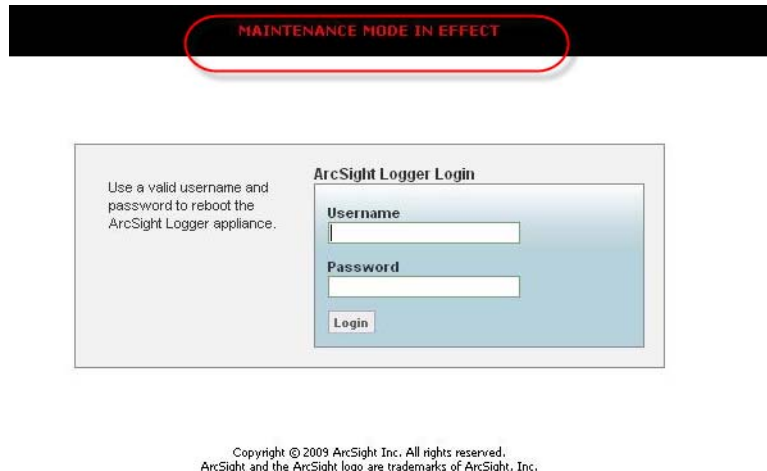
Logger users who will be performing operations that require it to be in maintenance mode must have the “Enable Maintenance Mode” privilege checked (**System Admin > User Management > Groups tab > System Admin Group**).



When a Logger is in maintenance mode, users with the “Enable Maintenance Mode” privilege can login but see this UI message:



All other users cannot login. The login screen displays this message:



Entering Maintenance Mode

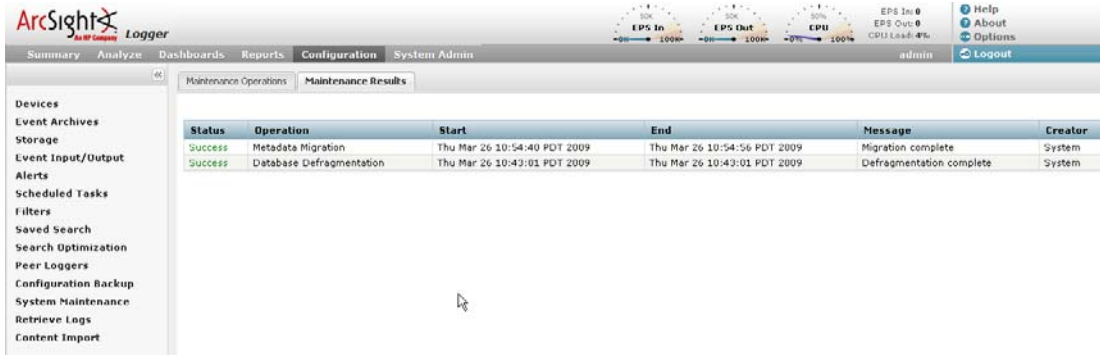
You cannot place a Logger in maintenance mode directly. A Logger can enter maintenance mode only when you perform an operation that requires it to be in that mode. For example, when defragmenting database, the user interface prompts you to enter Logger in maintenance mode, as illustrated in [“Database Defragmentation” on page 335](#).

Exiting Maintenance Mode

To exit maintenance mode, reboot the Logger appliance or restart the software Logger.

Checking Status of a Maintenance Operation

You can check the status of a maintenance operation on the Maintenance Results page. To access the Maintenance Results page (as shown in the example below), click **Configuration** (or **Configuration > Settings**) > **System Maintenance > Maintenance Results**.



Database Defragmentation

Logger's database can get fragmented over time. Frequent retention tasks can exacerbate this issue. The following symptoms are observed on a Logger when the database should be fragmented:

- Slow search and reporting
For example, even a search operation over the last two minutes of data is slow.
- Long pauses in the receiver and forwarder operations

You can defragment a Logger that exhibits the above listed symptoms. Make sure that you have read the following guidelines before starting the defragmentation process.

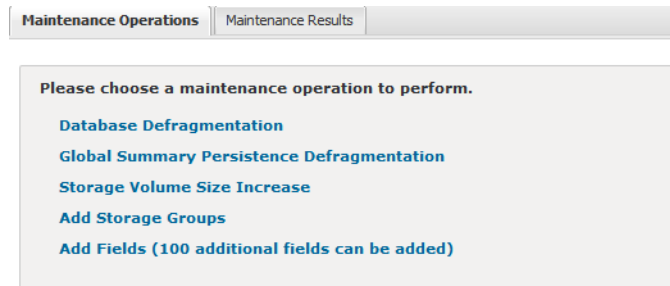
Guidelines for Database Defragmentation

- Ascertain that the Logger symptoms are not due to issues related to network infrastructure such as network latency or unexpected load on the Logger.
- The Logger system needs to be placed in maintenance mode before defragmentation can begin. As a result, most processes on the Logger are stopped—no events are processed or scheduled jobs run, and most user interface operations are unavailable. For more information about maintenance mode, see [“System Maintenance” on page 333](#).
- A minimum amount of free disk space is required on your system to run database defragmentation. The utility automatically checks for the required free space and displays a message if sufficient disk space is not found.
- Although you can defragment as needed, if you are using this utility too often (such as on a system that was defragmented over the last few days), contact customer support for guidance.
- If the defragmentation process fails at any point, the Logger returns to the same state that it was in before you started defragmentation.
 - ◆ You can safely reboot the Logger appliance and restart the process from the beginning.
 - ◆ For the software Logger, restart the Logger process as described in [“Process Status” on page 415](#).
- You can perform this process only if you have the “Enable Maintenance Mode” privilege set to Yes (**System Admin > User/Groups > Manage Groups > System Admin Group**).

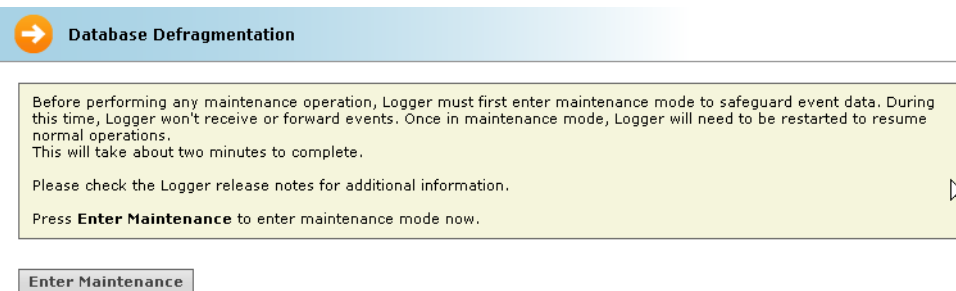
Defragmenting a Logger

To defragment a Logger:

- 1 Click **Configuration** (or **Configuration > Settings**) > **System Maintenance**. The Maintenance Operations panel displays the available options.



- 2 Click **Database Defragmentation**.
- 3 Click **Enter Maintenance** so that the Logger can enter maintenance mode. For more information about maintenance mode, see [“System Maintenance” on page 333](#).



- 4 A minimum amount of free storage is required for the database defragmentation process to proceed. Therefore, Logger performs a check to determine free storage when entering maintenance mode.
 - ◆ If the required storage is not found, follow the instructions found in [“Freeing storage space for defragmentation” on page 337](#).
 - ◆ If the required amount of free storage is found and Logger successfully enters maintenance mode, the following screen is displayed. Click **Begin Defragmentation** to start the defragmentation process.



Note

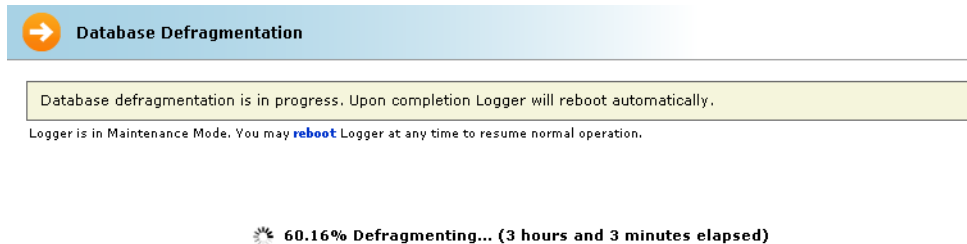
On the software Logger, the following Database Defragmentation screens instruct you to click **Restart** to resume normal operation when Logger is in maintenance mode. When you click restart, only the Logger service and its related processes are started on the machine on which the software Logger is installed.



Figure 6-17 Begin Database Defragmentation

- 5 The defragmentation process starts. A progress indicator shows the status of defragmentation, as shown in the example below. HP recommends that you do not attempt any operation on the Logger until defragmentation has completed.

Once defragmentation is complete, the Logger reboots automatically. This exits maintenance mode.



Freeing storage space for defragmentation

If the required storage is not found, Logger prompts you to free sufficient space, as shown in the following example:



Note

The Manual Deletion option (shown in the following figure) is not available on L7x00 Loggers.

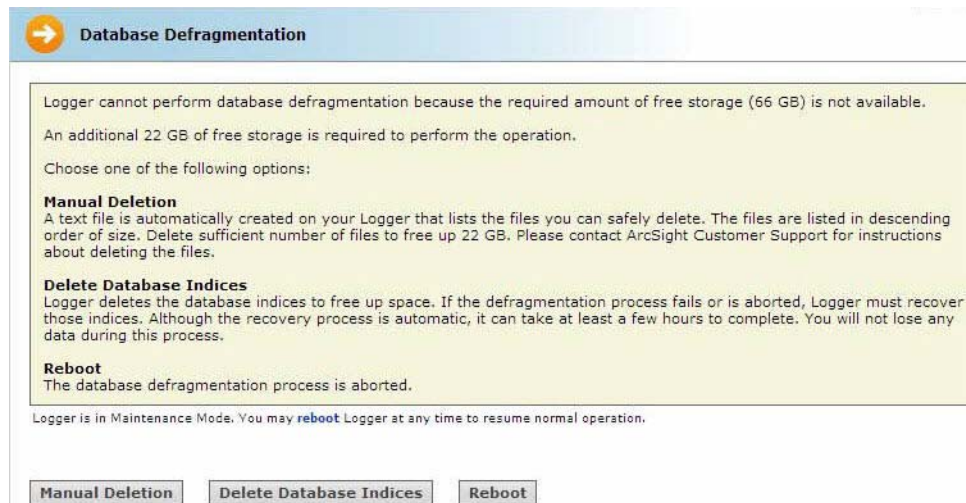
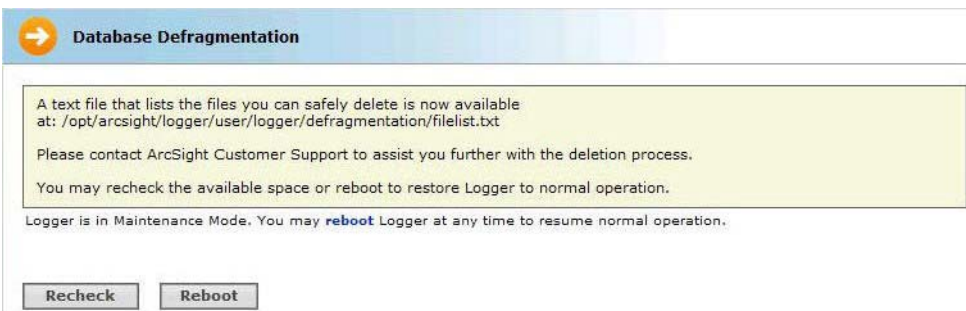


Figure 6-18 Required storage for Database Defragmentation is not available

You can choose from one of the following options:

◆ **Manual Deletion**

A text file is automatically created on your Logger that lists the files you can safely delete. The figure below is for a Logger appliance. On software Loggers, this file is located in `<install_dir>/current/arcsight/logger/user/logger/defragmentation/filelist.txt`.



The files are listed in descending order of size in the text file. You can delete sufficient number of files to free up storage. However, **do not** delete the files before contacting customer support for instructions and guidance.

Follow these steps to proceed:

- i Leave the message screen without taking any action.
- ii Contact customer support for instructions on deleting files listed in the text file.
- iii After deleting sufficient number of files, resume the Database Defragmentation process from the message screen in [Step i on page 338](#). To resume, click **Recheck** to check whether sufficient storage is now available for defragmentation to proceed.

If sufficient storage is found, the screen in [Figure 6-17 on page 337](#) is displayed. Click **Begin Defragmentation** to proceed further.

If sufficient storage is still not found, the screen in [Figure 6-18 on page 338](#) is displayed. Choose from the listed options to create additional space. See [“You can choose from one of the following options:” on page 338](#) for more information.



If you need to exit the defragmentation process without creating sufficient storage, click **Reboot**.

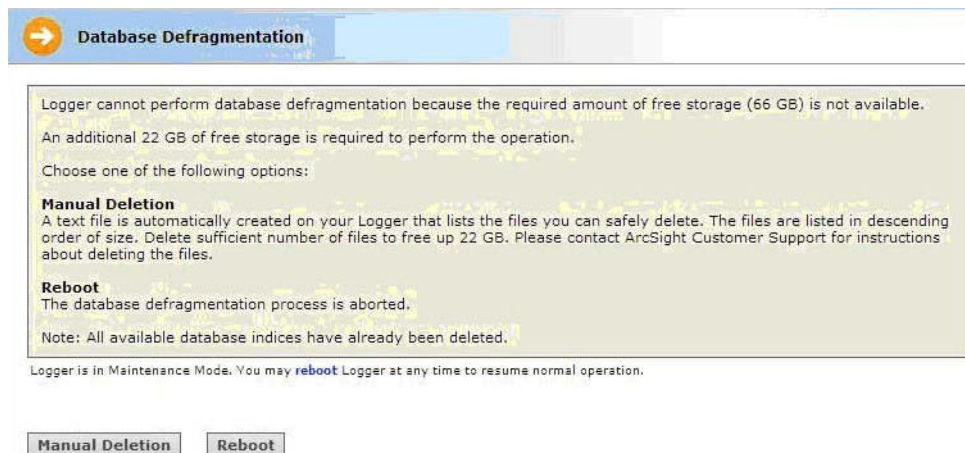
◆ Delete Database Indices

Logger automatically deletes a sufficient number of database indices, starting with the largest index, to free up the required amount of storage. If sufficient space becomes available after deleting database indices, defragmentation proceeds further automatically.

However, if sufficient storage is not available even after dropping database indices, the following screen is displayed.



The Manual Deletion option (shown in the following figure) is not available on L7x00 Loggers.



Follow these steps to proceed:

i Click **Manual Deletion**.

A text file is created on your Logger that lists the files you can safely delete. The files are listed in descending order of size in a text file.

ii Click **Reboot**.

Logger exits the maintenance mode.

iii Contact customer support for instructions on manually deleting the files.

You can delete sufficient number of files to free up storage.

- iv After deleting the files, restart the defragmentation process from [Step 1 on page 336](#).



Note

If the defragmentation process fails or is aborted at any time, Logger must recover those indices. Although the recovery process is automatic, it can take at least a few hours to complete. You will not lose any data during this process.

◆ **Reboot**

The database defragmentation process is aborted and Logger returns to the state it was in before you started the defragmentation utility.

Global Summary Persistence Defragmentation

There is a known issue with the Global Summary Persistence functionality in version 5.3 of Logger. This feature is designed to persist the statistics reported in the global summary section of Logger through a reboot. In some environments, disk space may be affected due to this feature.

This release turns off the Global Summary Persistence functionality. As soon as possible after upgrading to Logger 5.3 SP1 or later, enter System Maintenance mode and defragment the Global Summary table. Make sure that you have read the following guidelines before starting the defragmentation process.

Guidelines for Global Summary Persistence Defragmentation

- The Logger system needs to be placed in maintenance mode before Global Summary Persistence defragmentation can begin. As a result, most processes on the Logger are stopped—no events are processed or scheduled jobs run, and most user interface operations are unavailable. For more information about maintenance mode, see [“System Maintenance” on page 333](#).
- A minimum amount of free disk space is required on your system to run Global Summary Persistence defragmentation. The utility automatically checks for the required free space and displays a message if sufficient disk space is not found.
- If the defragmentation process fails at any point, the Logger returns to the same state that it was in before you started defragmentation. You can safely reboot the appliance or restart the software Logger process and try again.
 - ◆ Reboot the Logger appliance as described in [“System Reboot” on page 358](#).
 - ◆ For the software Logger, restart the Logger process as described in [“Process Status” on page 415](#).
- You can perform this process only if you have the “Enable Maintenance Mode” privilege set to Yes (**System Admin > User/Groups > Manage Groups > System Admin Group**).

To defragment for the Global Summary Persistence issue:

- 1 Click **Configuration** (or **Configuration > Settings**) > **System Maintenance**. The Maintenance Operations panel displays the available options.

Maintenance Operations | Maintenance Results

Please choose a maintenance operation to perform.

Database Defragmentation
Global Summary Persistence Defragmentation
Storage Volume Size Increase
Add Storage Groups
Add Fields (100 additional fields can be added)

2 Click **Global Summary Persistence Defragmentation**.
3 Click **Enter Maintenance** so that the Logger can enter maintenance mode. For more information about maintenance mode, see [“System Maintenance” on page 333](#). The Global Summary Persistence Panel displays information about the operation.

ArcSight[®]
An HP Company
Logger

50K
EPS In

50K
EPS Out

50%
CPU

EPS In: 0
EPS Out: 0
CPU Load: 1%

Help
About
Options

Summary | Analyze | Dashboards | Reports | Configuration | System Admin | admin | Logout

Global Summary Persistence Defragmentation

Logger is ready to perform the Global Summary Persistence Defragmentation.
This should take approximately 25 minutes.
Please check the logger release notes for additional information.
Click **Begin Global Summary Persistence Defragmentation** to begin.

Logger is in Maintenance Mode. You may **restart** Logger at any time to resume normal operation.

Begin Global Summary Persistence Defragmentation

Figure 6-19 Begin Global Summary Persistence Defragmentation

- 4 Click **Begin Global Summary Persistence Defragmentation** to start the defragmentation process.
- 5 The defragmentation process starts. A progress indicator shows the status of defragmentation. HP recommends that you do not attempt any operation on the Logger until defragmentation has completed.

Once defragmentation is complete, the Logger reboots or restarts. This automatically exits maintenance mode.

Note

On software Loggers, only the Logger service and its related processes are restarted.

Confidential

Logger Administrator's Guide **341**

Storage Volume Size Increase

You can extend the storage volume size you established during initialization at any time. Once extended, the volume size cannot be reduced. The Logger interface guides you about current and the maximum value to which you can increase the size.



For the “Storage Volume Size Increase” operation to show as an option under the System Maintenance operations (**Configuration** (or **Configuration > Settings**) > **System Maintenance**), you need to belong to the System Admin group (with “Enable Maintenance Mode” privilege enabled) and the Logger Rights group.

About Increasing Storage Volume Size on a SAN Logger

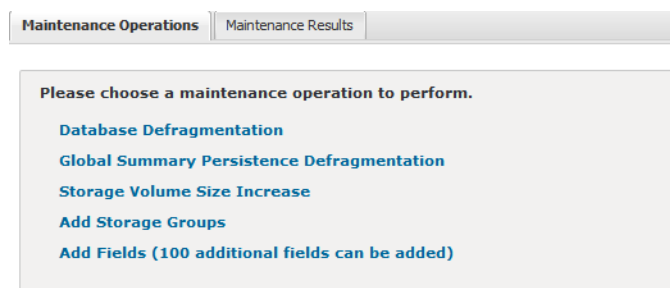
Logger cannot detect a resized LUN. Therefore, if you change the LUN size after it has been mounted on a Logger, the new size is not recognized by Logger. As a result, you can only increase the size of a storage volume to the LUN size that was initially mounted on the Logger. Currently, Logger supports up to a 5.4 TB LUN.

The following examples illustrate storage volume increase on a SAN Logger.

Initial LUN Size	LUN Resized	Current Storage Volume Size	Storage Volume Size Increase Allowed
4 TB	No	1 TB	Yes
4 TB	No	4 TB	No
4 TB	5 TB	1 TB	Yes, only up to 4 TB
2 TB	4 TB	1 TB	Yes, only up to 2 TB

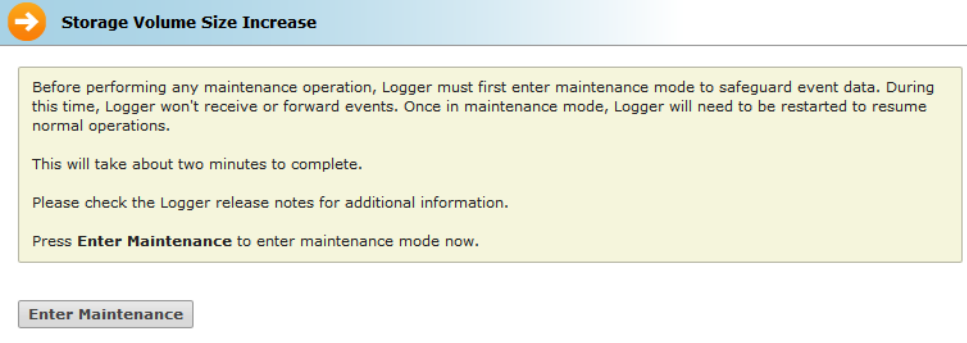
To Increase the size of a storage volume:

- 1 Click **Configuration** (or **Configuration > Settings**) > **System Maintenance**. The Maintenance Operations panel displays the available options.



- 2 Click **Storage Volume Size Increase**.
- 3 Click **Enter Maintenance** so that the Logger can enter maintenance mode.

For more information about maintenance mode, see [“System Maintenance” on page 333](#).



Storage Volume Size Increase

Before performing any maintenance operation, Logger must first enter maintenance mode to safeguard event data. During this time, Logger won't receive or forward events. Once in maintenance mode, Logger will need to be restarted to resume normal operations.

This will take about two minutes to complete.

Please check the Logger release notes for additional information.

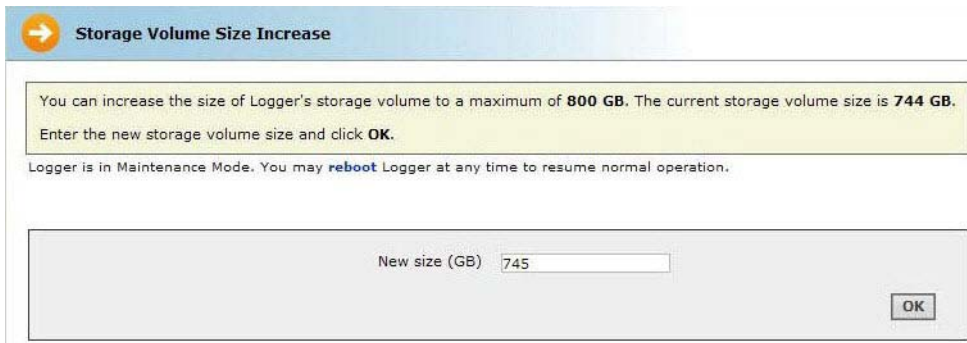
Press **Enter Maintenance** to enter maintenance mode now.

Enter Maintenance

- 4 While entering the maintenance mode, Logger performs a check to determine if the storage volume size can be increased and by what amount. If the storage volume can be increased, a message similar to the following is displayed. Enter the new size and click **OK**.



On the software Logger, the following Storage Volume Size Increase screens instruct you to click **restart** to resume normal operation when Logger is in maintenance mode. When you click restart, only the Logger service and its related processes are restarted.



Storage Volume Size Increase

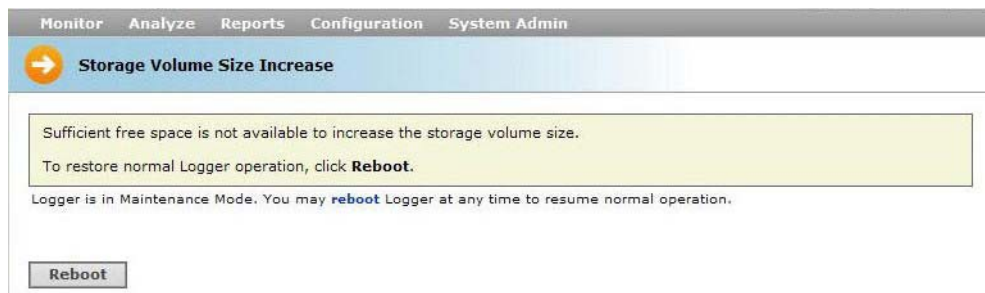
You can increase the size of Logger's storage volume to a maximum of **800 GB**. The current storage volume size is **744 GB**. Enter the new storage volume size and click **OK**.

Logger is in Maintenance Mode. You may **reboot** Logger at any time to resume normal operation.

New size (GB)

OK

If sufficient space is not found to increase the storage volume, the following message is displayed. Click **Reboot** to restart the Logger and exit the maintenance mode.



Storage Volume Size Increase

Sufficient free space is not available to increase the storage volume size.

To restore normal Logger operation, click **Reboot**.

Logger is in Maintenance Mode. You may **reboot** Logger at any time to resume normal operation.

Reboot

Adding Storage Groups

In addition to the two storage groups that exist on your Logger by default, you can add up to four additional storage groups. Prior to Logger 5.2, the additional storage groups had to be added at Logger initialization time. You could not add them later, once Logger was initialized. Starting with Logger 5.2, you can add storage groups at any time if the following conditions are met:

- The maximum allowed six storage groups do not exist on your Logger already.
- The storage volume contains spare storage space that can be allocated to the storage groups you will add.



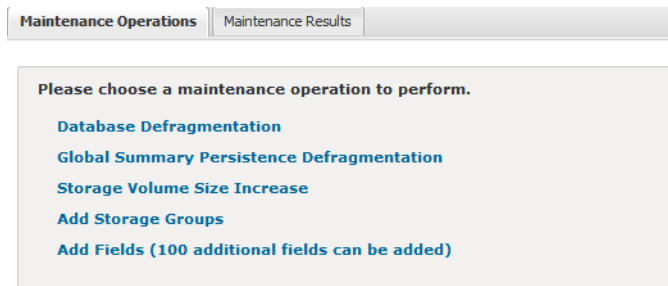
If you do not have sufficient space in the storage volume to add another storage group and the existing groups have free space, consider reducing the size of existing storage groups to make space available for the storage groups you want to add. Alternatively, increase the size of your existing storage volume, as described in [“Storage Volume Size Increase” on page 342](#).

The Logger must be in maintenance mode when adding storage groups. When you add a storage group, Logger automatically checks to ensure that the storage group size you specified is greater than the minimum size required (5 GB) and less than the amount of space available in the storage volume.

Once you have added storage groups and rebooted your Logger to exit the maintenance mode, remember to configure the Archive Storage Settings for the groups you just added so that event archives are created for them.

To add a storage group:

- 1 Click **Configuration** (or **Configuration > Settings**) > **System Maintenance**. The Maintenance Operations panel displays the available options.



- 2 Click **Add Storage Groups**.

A maximum of six storage groups can exist on Logger. Therefore, you can add up to four storage groups in addition to the two that exist by default on Logger.

If the maximum number of allowed storage groups **do not** exist on Logger, a screen prompts you to enter maintenance mode, as described in the next step.

If all six storage groups exist on Logger or sufficient space does not exist in the storage volume to add additional group, a message is displayed on your screen and the Logger cannot enter maintenance mode.

- 3 Click **Enter Maintenance** so that the Logger can enter maintenance mode.

For more information about maintenance mode, see [“System Maintenance” on page 333](#).

- 4 Once Logger enters maintenance mode, the following Add Storage Groups page is displayed.

→ Add Storage Groups

The storage volume has 40 GB left of unallocated space.

Logger is in Maintenance Mode. You may **restart** Logger at any time to resume normal operation.

Name	Maximum Age (Days)	Maximum Size (GB)	Creator	Last Editor
Default Storage Group	180	30	admin	admin
Internal Event Storage Group	365	30	System	System

Name

Maximum Age (Days)

Maximum Size (GB)

Add

This screen also lists information about the existing storage groups and the amount of space remaining in the storage volume.

- Enter the following information.

Parameter	Description
Name	Choose a name for the storage group
Maximum Age (Days)	Specify the number of days to retain events. Events older than this number of days are deleted.
Maximum Size (GB)	Enter a maximum event data size, in GB.

- Click **Add**.

The storage group is added to your Logger. If your Logger has not reached the maximum allowed six storage groups, you can click **Add** to add more storage groups. However, if the maximum number has been reached, the Add button is not displayed. If you do not want to add more storage group, go to the next step.

- Reboot your Logger appliance or restart software Logger for changes to take effect and for the appliance to exit the maintenance mode

Adding or Importing Schema Fields

The Logger schema contains a predefined set of fields. A field-based query can contain only these fields. Additionally, you can index only these fields for faster search operations. For instructions on how to view the default Logger Schema fields, see [“Viewing Default Fields” on page 324](#).

Prior to Logger 5.2, if your log analysis needs required you to search on a field that is currently not present in the Logger schema, you did not have a way of adding it to the schema yourself. Starting with Logger 5.2, you can add additional fields to the Logger schema. That is, you can insert fields in your Logger schema that are relevant to the events you collect on your Logger, thus enabling you to search and report using these fields.

Additionally, you can index the fields you add so that the search and report queries that use these fields run faster. For example, a financial institution might want to add credit card numbers or social security numbers to the schema.

You can add up to 100 custom schema fields on Logger. You can also import custom fields from a peer Logger. However, the total number of added and imported fields cannot exceed the maximum allowed 100 fields.

You can index up to 123 fields on Logger. Therefore, the number of custom schema fields you can index will depend on the number of default fields you currently have indexed on your Logger.

The events that contain custom fields must be in CEF format (key-value pairs) for Logger to process them. Therefore, you will need to either use a SmartConnector that generates additional data or define an ArcSight FlexConnector to collect and parse events containing custom fields from the event source, convert them into CEF format, and forward them to the Logger.

Logger can only process events from FlexConnectors written using connector build 5.0.0.5560 or later. For details about designing FlexConnectors, see the ArcSight FlexConnector Developer's Guide.

**Note**

Logger cannot process the additional fields data received in CEF version 0 from a FlexConnector, and assumes a NULL value for such fields when they are present in a CEF version 0 event. As a result, you cannot search on these fields or index them. However, these fields are displayed in the UI display when you select "*" in the fieldset because the interface displays information contained in the raw event. Therefore, if Logger receives "ad.callnumber=5678", the Logger UI will display a column, ad.callnumber, with value 5678. However, a search on "5678" will not return this event in the search results.

You need to be in maintenance mode to add or import custom schema fields. The process of adding or importing schema fields involves an add or import operation followed by a save operation. The add or import operation adds the specified fields but does not write them to the Logger schema. You can edit or delete the added or imported fields at this point. Once you save these fields, the fields are written to the schema. From this point on, these fields cannot be edited or deleted. Therefore, carefully review the fields you are adding to the schema before saving them.

**Note**

For the "Add Fields" operation to show as an option under the System Maintenance operations (**Configuration** (or **Configuration > Settings**) > **System Maintenance**), you need to belong to the System Admin group (with "Enable Maintenance Mode" privilege enabled) and the Logger Rights group.

You need to specify the following information to add a custom schema field:

- **Display name**

A meaningful name for the field. This name is displayed as the column header name for the field and is the one you specify in a search query. For example, SocialSecurityNumber.

■ Type

The type of data this field will contain. The available options are Double, BigInt, DateTime, Text.

The following table describes each data type.

Type	Description
Double	Use to store decimal numbers or fractions. Numbers from -1.79769313486231570E+308 through -4.94065645841246544E-324 for negative values and 4.94065645841246544E-324 through 1.79769313486231570E+308 for positive values.
BigInt	Use to store whole numbers. Numbers from -2^{63} (-9,223,372,036,854,775,808) through $2^{63}-1$ (9,223,372,036,854,775,807)
DateTime	Use to store both dates and time or only dates.
Text	Use to store any characters. You can store a maximum of 255 characters per field.

■ Length

This field is only relevant when the Type specified is Text. This field specifies the maximum number of characters allowed in the value of the field when the data type is Text.

■ Field name

The field name that you want to add to the Logger schema. Typically, this is an abbreviated version of the Display name. For example, SSN.

Importing Schema Fields from Peers

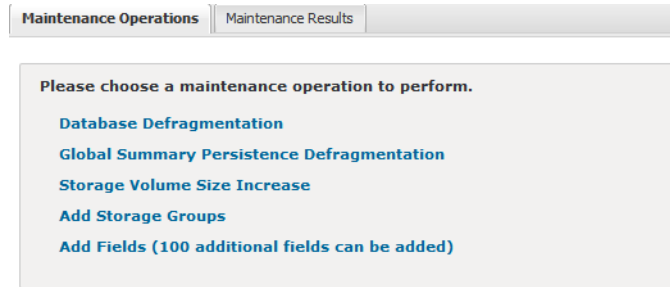
If your Logger is a peer of another Logger, you can import the custom fields added to the peer's schema. You specify the peer from which you want to import fields in the user interface screen. Fields can be imported if the following conditions are met:

- A field of the same Display name and Field name does not exist on the Logger to which you are importing schema fields. If conflicting fields exist, they are still imported but are flagged in the user interface screen. You cannot save the imported fields to schema until you resolve the conflicts.
- A maximum of 100 custom fields has not been reached on the importing Logger. If there are more fields than can be imported, only the first *N* until the allowed maximum is reached will be imported.

The custom schema fields contained in a search query must exist on all peers on which the query is run. Otherwise, the query will not run and return an error.

To add or import custom schema fields:

- 1 Click **Configuration** (or **Configuration > Settings**) > **System Maintenance**. The Maintenance Operations panel displays the available options.



- 2 Click **Add Fields (100 additional fields can be added)**.

You can add a maximum of 100 custom fields to Logger schema. The number in the "Add Fields" link reflects the number of custom fields you can add. This number decreases as you add fields to Logger schema.

- 3 Click **Enter Maintenance** so that the Logger can enter maintenance mode.

For more information about maintenance mode, see ["System Maintenance" on page 333](#).

- 4 Once Logger enters maintenance mode, the Add Fields page is displayed.

You can add fields manually or import them from a peer Logger.

To manually add fields:

- 1 After entering Maintenance Mode, click "Add a New Field", if it is not selected.

- 2 Enter a meaningful name in the Display Name field.

This name is the one you specify in a search query and is displayed as the column header name for the field in search results. For example, SocialSecurityNumber. This name is not added to the Logger schema. Follow these guidelines when specifying a display name:

- ◆ The name can contain up to 100 characters.
- ◆ The name can contain alphanumeric characters, hyphens ("-"), and underscores ("_"). However, a hyphen ("-") or an underscore ("_") cannot be the first character in the name. Additionally, the name cannot begin with "arc_".
- ◆ The name must be unique; that is, another field (custom or Logger schema) of the same display name must not already exist on the Logger.
- ◆ Only ASCII characters are allowed. That is, no native Chinese or Japanese characters are accepted in this field.

- 3 Select a data type for the field from the Type drop-down menu.

- 4 The available options are Double, BigInt, DateTime, Text. See ["Type" on page 347](#) for more information.

- 5 In the Length field, enter the maximum number of characters allowed in the value of the field *when the data type is Text*. This field is only available when the Type specified is Text. You can specify from 1 to 255 characters in this field.

- 6 Enter a name in the Field name field.

This is the name that will be added to the Logger schema. Typically, this is an abbreviated version of the Display name. For example, SSN. Follow these guidelines when specifying a Field name:

- ◆ This is a required field.

- ◆ The name can contain up to 40 characters and can contain alphanumeric, hyphen ("-"), and underscore ("_") characters. Underscore ("_") is used as an escape character for the actual field name. Therefore, the underscore ("_") you specify in the field name is converted to a double underscore ("__") in the actual field name.
- ◆ The name must be unique; that is, a custom field of the same Field name must not already exist on the Logger.
- ◆ Only ASCII characters are allowed. That is, no native Chinese or Japanese characters are accepted in this field.

Once you enter a name in this field, a prefix and a suffix is automatically added to it, and the resulting name is displayed in the Actual Field Name field, as shown in the following figure. This field displays the way the field name you entered earlier will be stored on Logger. The prefix, "ad." signifies "additional data" and the suffix signifies the data type of the field. The Actual Field Name field is a non-editable field and is displayed on the user interface only for your reference.

7 Click **OK**.

The field you added is displayed in the upper section of the Add Fields form, as shown in the following figure. This field is not saved yet (in "Ready to Save" state) and you can edit or delete it. Once you click Save, the field is added to the schema and cannot be changed or deleted.

Logger is ready for adding new fields. You can add up to **96** additional fields.
The fields in "Ready to save" status are not in logger schema yet. Click Save to write these fields to the schema.

Logger is in Maintenance Mode. You may **restart** Logger at any time to resume normal operation.

Save



Display Name	Type	Length	Field Name	Actual Field Name	Creator	Created	Status
testBigInt	BIGINT	-	BigIntField	ad.BigIntField.i	admin	Jan 14, 2014 3:09:32 PM PST	Ready for save
testDateTime	DATETIME	-	DateTimeField	ad.DateTimeField.d	admin	Jan 14, 2014 3:09:45 PM PST	Ready for save
testDouble	DOUBLE	-	DoubleField	ad.DoubleField.r	admin	Jan 14, 2014 3:09:15 PM PST	Ready for save
testText	TEXT	255	TextField	ad.TextField	admin	Jan 14, 2014 3:08:52 PM PST	Ready for save

You can import fields from peer Loggers. Make sure this Logger is configured as the peer of the Logger from which you want to import fields.

Display Name:
 Type: **DOUBLE**
 Field Name:

OK

8 Follow [Step 1](#) through [Step 7](#) to add additional fields.

- 9 Review the added fields and make any edits () or deletions (), if necessary.



- The next step commits the added fields to Logger's schema. This process is irreversible; that is, once the fields are written to Logger's schema, they cannot be edited or deleted.
- If you exit this process without saving, the fields you were adding are not remembered and your changes are lost.

- 10 Click **Save** to commit the added fields and write them to your Logger's schema.



To import fields from a peer:

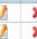



- 1 Click "Import Fields From Peers".
- 2 Select the peer from which you want to import the fields from the Peer Host Name drop-down list.

- 3 Click **OK** in the bottom right corner of the screen.

If there are no conflicting fields, all fields from the peer are imported successfully.

If there are conflicts, the conflicting fields are displayed ahead of the ones that were imported successfully. The Status column describes the reason for the conflict. You must fix the listed issues before you can save these fields to the schema. Use the edit

() or delete () icon to make changes or delete the added fields.

Display Name	Type	Length	CEF Name	Creator	Created	Status		
sjc_char2	TEXT	20	ad.char2	admin	Jul 12, 2011 2:24:58 AM PDT	Another field of the same display name, 'sjc_char2', exists. Enter another display name.		
sjc_char4	TEXT	20	ad.char4	admin	Jul 12, 2011 2:24:58 AM PDT	Ready for save		
sjc_char2	TEXT	20	ad.char2	admin	Jul 12, 2011 2:24:39 AM PDT	Saved		
sjc_char3	TEXT	20	ad.char3	admin	Jul 12, 2011 2:24:39 AM PDT	Saved		

If there are more fields than can be imported, only the first N until the allowed maximum (100) is reached will be imported.



The imported fields are not committed to Logger's schema yet. The next step commits them. This process is irreversible; that is, once the fields are written to Logger's schema, they cannot be edited or deleted.

If you exit this process without saving, the fields you were adding are not remembered and your changes are lost.

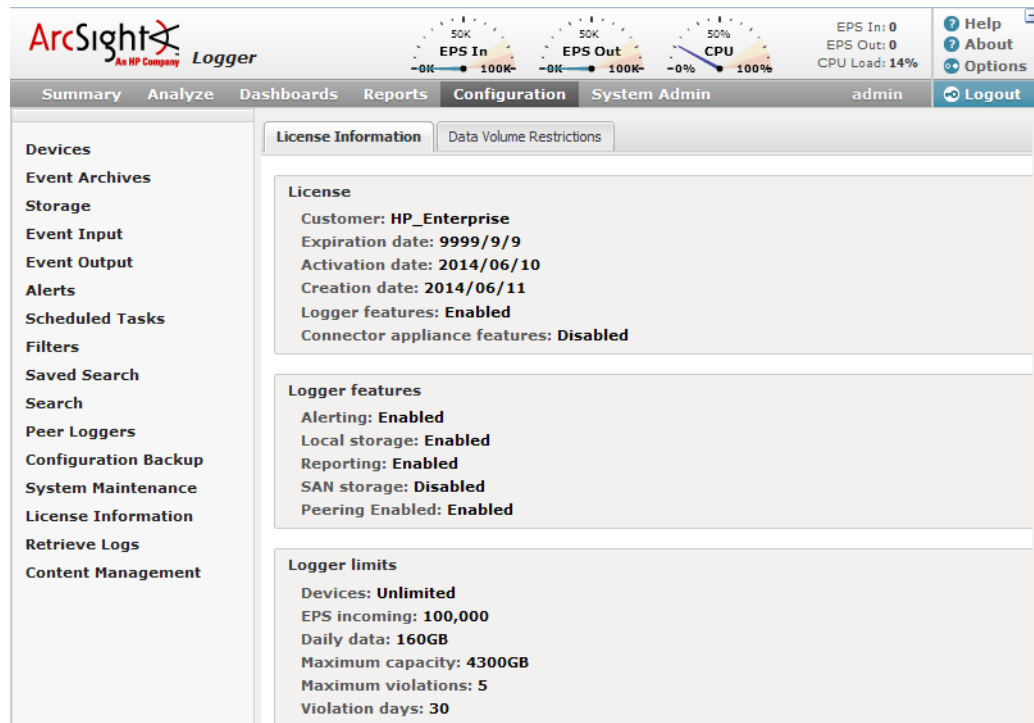
- 4 Click **Save** to commit the added fields and write them to your Logger's schema.

To view existing custom schema fields:

See [“Viewing Custom Fields” on page 324](#).

License Information

The License Information page (**Configuration** (or **Configuration > Settings**) > **License Information**) provides information about the currently applied license, as shown in the following figure.



To upload a new license, open **System Admin** from the top-level menu bar, and then click **License & Update** in the **System** section. For details, see [“License & Update” on page 364](#) for Logger appliances or [“License & Update” on page 414](#) for software Loggers.

Data Volume Restrictions

The Data Volume Restrictions page lists the data stored on your software version of Logger on day-by-day basis in the last 30 days. It also indicates the days on which data limits were exceeded, as shown in the following figure. If the data-limit has been exceeded 6 times, you cannot search on Logger system and need to wait until the listed 30 days have 5 or less violations.

Summary Analyze Reports Configuration System Admin admin Logout			
Devices Event Archives Storage Event Input Event Output Alerts Scheduled Tasks Filters Saved Search Search Peer Loggers Configuration Backup System Maintenance License Information Retrieve Logs Content Management	License Information		Data Volume Restrictions
	Date	Daily Data (MB)	Limit Exceeded
	12/14/13	0	false
	12/15/13	0	false
	12/16/13	0	false
	12/17/13	0	false
	12/18/13	0	false
	12/19/13	0	false
	12/20/13	0	false
	12/21/13	0	false
	12/22/13	0	false
	12/23/13	0	false
	12/24/13	0	false
	12/25/13	0	false
	12/26/13	0	false
	12/27/13	0	false
	12/28/13	0	false
	12/29/13	33844	true
	12/30/13	0	false
	12/31/13	0	false

Retrieve Logs

Logger records some audit and debug information, including details of any issues that occur. These system logs (not be confused with the event logs that Logger was designed to process), are like the “black box” on an airliner. If something goes wrong, the logs can be helpful.

Customer support may ask you to retrieve logs as part of an incident investigation. If so, follow the steps below and upload the resulting .zip file to customer support.

To retrieve Logger system logs:

- 1 Click the **Configuration** (or **Configuration > Settings**) > **Retrieve Logs**.

The page shown in [Figure 6-20](#) appears.

- 2 When the Summary Status is Completed, click **Download** to retrieve the system log files are compressed into a single zip file.


Retrieve Snapshot Status		
Summary		
Name	Thread-42	
Request ID	5ekaeDsBABCAA_dhKloGMw	
Processing Time	40 sec 50 ms	
Status	 Processing...	
Action	Start Time	Time to Complete
Database content	12/7/12 5:22 PM	1 sec 721 ms
Report content	12/7/12 5:22 PM	1 sec 166 ms
Retrieving logs	12/7/12 5:22 PM	Processing...
Download		

Figure 6-20 Retrieve Logs provides snapshot status.

Content Management

Depending on their rights, users can export Alerts, Dashboards, Filters, Parsers, Saved Searches, and Source Types from a Logger to a file, and then import that content onto another Logger or re-import it onto that same Logger, as a backup. For information on the user rights necessary to import or export a particular type a content, instructions, and guidelines for importing and exporting Logger content, see [“Exporting Content” on page 354](#) and [“Importing Content” on page 353](#).

Content import and export is useful in these situations:

- When you want to make a backup of Logger content. If your Logger becomes unavailable or is reset to its factory defaults, you can quickly restore its content by importing the saved content.
- When multiple Loggers with the same content need to be installed in your network, you need to configure only one Logger. Subsequent Loggers can be deployed by importing the first Logger’s content on them, thus reducing deployment time.
- When you want to add content from one Logger to the content on another.

Using the Export function, you save the content from a Logger to a storage location on your network or to the local disk of the computer from which you connect to the Logger. When you need to use that content for any of the situations described previously, simply import the saved content.

Importing Content

The content you are able to import depends on your user rights. If you have any of the following rights, the Import tab displayed in [Figure 6-21](#) is available:

- Logger **Rights > Filters**: Edit, save, and remove shared filters.
- Logger **Rights > Forwarders and Alerts**: Edit, save, and remove forwarders and alerts.



While this Logger right enables you to edit, save, and remove both forwarders and alerts, you can only import alerts, but not forwarders.

- Logger **Rights > Dashboards**: Edit, save, and remove dashboards.
If the user has the dashboard save right but does not have the saved search save right, then the dashboards using search results panels will not be imported (A warning message will indicate which dashboards are skipped).
- Logger **Rights > Saved Search**: Edit, save, and remove saved search.
- **System Admin**: For parsers and source types, the user can be assigned to any System Admin Group. If the user is not an admin, then Parsers and Source Types are not importable.

Even if you see the import tab, you may not be able to import all of the content types. If you do not have one of the above user rights, then you cannot import that type of content and will get a warning message instead.

Importing Guidelines

Make sure you are familiar with these guidelines before importing Logger content:

- If an object with the same name exists on the importing system, object being imported is named *<ObjectName> [import]*. For example, an imported alert is named *AlertName [import]* and an imported filter is named *FilterName [import]*.

If an object with the name *<ObjectName> [import]* already exists on the importing Logger (from a previous import procedure), the object being imported is named *<ObjectName> [import] [import]*.

- Be sure to set the alert destinations (SNMP, Syslog, ESM Destination, and SMTP servers) for alerts you import because this information is not included in the exported content.

To Import content from another Logger:

- 1 Click **Configuration** (or **Configuration > Settings**) on the top-level menu bar.
- 2 Click **Content Management** in the left panel.
- 3 Click the **Import** tab.

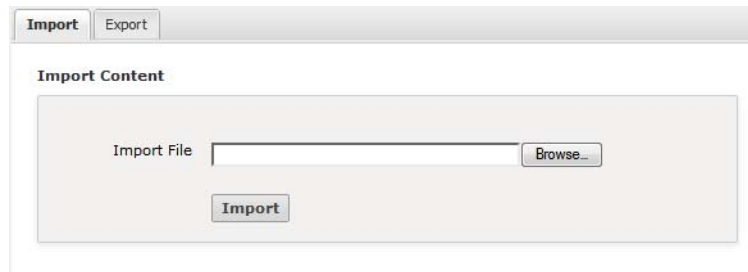


Figure 6-21 Import tab

- 4 Click **Browse** to locate the file

The file must reside on a local or remote drive accessible to the system whose browser you are using to access Logger's user interface.

- 5 Click **Import**.

Exporting Content

The content you are able to export depends on your user rights. If you have any of the following rights, the Export tab displayed in [Figure 6-22](#) is available:

- Logger **Rights > Filters**: Use and view shared filters.
- Logger **Rights > Forwarders and Alerts**: View forwarders and alerts.



Note

While this Logger right enables you to view both forwarders and alerts, you can only export alerts, but not forwarders.

- Logger **Rights > Dashboards**: Use and view dashboards.

If the user has the dashboard read right but does not have the saved search read right, then dashboards having search results panels are not available for selection from the Content to Export dialog box.

- **Logger Rights > Saved Search:** View saved search.
- **System Admin:** For parsers and source types, the user can be assigned to any System Admin Group. If the user is not an admin, then Parsers and Source Types are not exportable.

Even if you see the Export tab, you may not be able to export all of the content types. If you do not have one of the above user rights, then the corresponding type of content type is not available in the Content to Export dialog box.

Exporting Guidelines

Make sure you are familiar with these guidelines before exporting Logger content:

- The exported content is in XML format in a gzip file. For example, allfilters.xml.gz.
- The folder on the remote file system to which you are exporting Logger content needs to exist before you can export content to it.
- When exporting alerts, the query associated with the alert, match count, threshold, and status are included in the export. The export does not include e-mail, SNMP, ESM Destination information, or syslog destination information. Since alert destination (SNMP, Syslog, ESM Destination, and SMTP servers) information is not exported, you will need to set this information for alerts you import.
- When exporting dashboards, the content of any saved searches used in the exported dashboards is also exported.
- When exporting source types, the content of the parsers used in the exported source types is also exported.

To export Logger content:

- 1 Click **Configuration** (or **Configuration > Settings**) on the top-level menu bar.
- 2 Click **Content Management** in the left panel.
- 3 Click the **Export** tab.

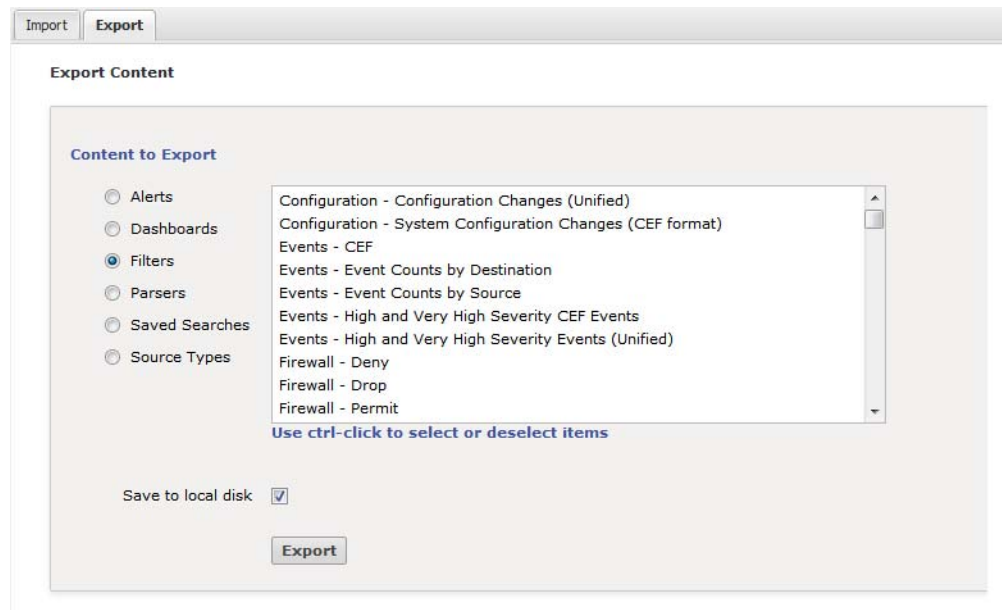


Figure 6-22 Export tab, with content type Filters selected

- 4 Click the button to select the type of content that you want to export. The displayed content changes with the type of content you select.
- 5 Select the objects to export in the **Choose Items to Export** field.
To select one object, click its name. To select multiple objects, hold the Ctrl key down and click the names.
- 6 For appliance Loggers, choose where to save the exported content. **Save to local disk** is the default option.

To save on the local disk of the computer from which you connect to the Logger, leave **Save to local disk** checked.

To export to a remote location:

- a Uncheck **Save to local disk** to display options for exporting to a remote file system.

The image shows two screenshots of the 'Content to Export' configuration window. The left screenshot shows the 'Save to local disk' checkbox checked and circled in red. The right screenshot shows the 'Export to remote file system' section with fields for 'Mount Location', 'Remote file path and name', and 'Overwrite if file exists', all highlighted with a red rounded rectangle.

- b Select the location to which you want to export the content in the Mount Location field. If the location you want is not in the drop-down list, you need to add it. For information about adding a network storage location, see ["Storage" on page 372](#).
- c In the **Remote file path and name** field, enter the folder location in which the exported contents file will be created at the Mount Location you specified in the previous step. The folder location you specify in this step must already exist on the Mount Location. It is not created by the Logger.



Note

Specify the filename without using an extension.

- 7 Click **Overwrite if file exists** if you want to overwrite a file with the same name as the exported contents file in the folder location that you specified in the previous step.
- 8 Click **Export**.

System Admin - Logger Appliance

This chapter describes the System Administration tools that enable you to create and manage users and user groups, and to configure network, storage, and security settings for your system.

This chapter includes information on the following areas of system administration for the Logger Appliance:

- ["System Locale" on page 358](#)
- ["System Reboot" on page 358](#)
- ["Network" on page 358](#)
- ["SMTP" on page 363](#)
- ["License & Update" on page 364](#)
- ["Process Status" on page 364](#)
- ["SNMP" on page 365](#)
- ["SSH Access to the Appliance" on page 370](#)
- ["Audit Logs" on page 371](#)
- ["Audit Forwarding" on page 371](#)
- ["Remote File Systems" on page 372](#)
- ["SAN" on page 375](#)
- ["RAID Controller/Hard Disk SMART Data" on page 379](#)
- ["SSL Server Certificate" on page 381](#)
- ["SSL Client Authentication" on page 386](#)
- ["FIPS 140-2" on page 388](#)
- ["Authentication" on page 392](#)
- ["Login Banner" on page 401](#)
- ["User Management" on page 402](#)
- ["Change Password" on page 407](#)
- ["Monitoring System Health" on page 407](#)
- ["Using the Command Line Interface" on page 410](#)

System

From the System tab, you can configure system specific settings such as network settings (if applicable) and SMTP.

System Locale

The System Locale setting ensures that the user interface displays information such as date, time, numbers, and messages in the format and language appropriate for the selected country.

The System Locale is configured during the Logger installation process. Once configured it cannot be changed.

To view the System Locale:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **System Locale** in the **System** section.
The System Locale Setting dialog box displays the Locale.

System Reboot

To reboot or shutdown your system:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **System Reboot** in the **System** section.
- 3 Select from the following options:

Button	Description
Reboot	Your system reboots in about 60 seconds. The reboot process normally takes 5-10 minutes, during which time the system is unavailable.
Reboot in 5 Minutes	Your system reboots after a 5-minute delay. The reboot process normally takes 5-10 minutes, during which time the system is unavailable.
Shutdown	Automatically shuts down (powers off) the system.



Note

Each of the above actions can be cancelled. "Reboot" and "Shutdown" allow for cancellation within **60 seconds**. "Reboot in 5 Minutes" can be cancelled within **300 seconds**.

- 4 Click **Reboot**, **Reboot in 5 Minutes**, or **Shutdown** to execute the chosen action.



Caution

During reboot, Logger is not able to receive events. Events may be lost while Logger reboots, unless SmartConnectors are used. SmartConnectors cache events when destinations like Logger are temporarily unavailable.

Network



You can configure the DNS, Hosts, NICs, static routes, and system time settings from the Network menu.

System DNS

The **System DNS** tab allows you to edit the DNS settings and to add DNS search domains.

To change DNS settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** in the **System** section.
- 3 In the **System DNS** tab, enter new values for the IP address of the primary and secondary DNS servers, or edit the list of search domains.

To add a new domain, click the  icon. To remove a domain, click the  icon. To change the search order of domains, select a domain name, and click the up or down arrow until the domain is in the desired position.

- 4 Click **Save**.
- 5 Click **Restart Network Service** to put the changes into effect.

Hosts

The **Hosts** tab allows direct editing of your system's `/etc/hosts` file. You can enter data in the System Hosts text box or import it from a local file.

To change the Hosts information:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** in the **System** section, and then click the **Hosts** tab.
- 3 In the **System Hosts** text box, enter hosts information (one host per line) in this format:

```
<IP Address> <hostname1> <hostname2> <hostname3>
```

To import information from a file, click **Import from Local File**, and locate the text file on the computer from which you are accessing your system.

- 4 Click **Save**.

NICs

The **NICs** tab enables you to set the IP addresses for the network interface cards (NICs) on your system. Additionally, you can configure the hostname and default gateway for your system.

To set or change the NICs settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** in the **System** section.
- 3 In the **NICs** tab, enter the following settings. To edit the IP address, subnet mask, or speed/duplex of an NIC, select the NIC and click **Edit** above the NIC Name list.

Setting	Description
Default Gateway	The IP address of the default gateway.

Setting	Description
Hostname	<p>The network host name for this system. Make sure that your DNS can resolve the host name you specify to your system's IP address. Performance is significantly affected if DNS cannot resolve the host name.</p> <p>This name must be identical to the domain specified in the Certificate Signing Request, described in "Generating a Certificate Signing Request (CSR)" on page 383.</p> <p>Note: If you previously used a self-signed or CA-signed certificate on this system and are now changing its host name, you must regenerate a new self-signed certificate or CSR. Once obtained, the new certificate should be uploaded to ensure that the connectors (in FIPS mode) which communicate with your system are able to validate the host name. For more information about generating a CSR, see "Generating a Certificate Signing Request (CSR)" on page 383.</p>
Automatically route outbound packets (interface homing)	<p>When this option is enabled (checked box), the response packets are sent back on the same system interface on which the request packets had arrived. Enabling this option can improve performance as the routing decisions do not need to be made (using the default gateway information and static routes) to send packets out from your system. If you have static routes configured, they are ignored when this feature is enabled.</p> <p>When this feature is disabled (unchecked box), the static routes (if configured) are used to determine the interface through which the response packets should leave your system.</p> <p>If you configure only one network interface, this setting does not provide any additional benefit.</p>
IP Address	<p>The IP address for each network interface card (NICs) in your system. These IP addresses should be on separate subnets to avoid confusion and to allow load balancing between receivers and forwarders.</p> <p>Add NIC Alias</p> <p>You can create an alias for any listed NIC. To do so:</p> <ol style="list-style-type: none"> 1 Highlight the NIC for which you want to create an alias. 2 Click Add. 3 Create an alternative IP address for the alias. 4 Click Save. <p>You can identify the alias from its original by an appended colon alongside a digit indicating the number of aliases you have created on a particular NIC.</p> <p>Notes:</p> <ul style="list-style-type: none"> • You cannot alter the speed of an IP alias. • You can create as many aliases as you choose.
Subnet Mask	The subnet mask associated with the IP address you entered for an NIC.
Speed/Duplex	<p>Choose a speed and duplex mode, or let your system determine the network speed automatically:</p> <p>Auto (recommended)</p> <p>10 Mbps - Half Duplex</p> <p>10 Mbps - Full Duplex</p> <p>100 Mbps - Half Duplex</p> <p>100 Mbps - Full Duplex</p> <p>1 Gbps - Full Duplex</p>

4 Click **Save**.

- 5 Click **Restart Network Service** to put the changes into effect.

Static Routes

You can specify static routes for the NICs on your system.

To add, edit, or delete a static route:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** in the **System** section.
- 3 In the **Static Routes** tab:
 - ◆ To add a new static route, click **Add**.
 - ◆ To edit or delete an existing route, select the route first, then click **Edit** or **Delete**.

When adding or editing a static route, you need to configure these settings.

Setting	Description
Type	Whether the static route is to a Network or a Host
Destination	The IP address for the static route destination
Subnet Mask	The subnet mask if you specify a network as the destination
Gateway	The IP address of the gateway for the route

- 4 Click **Save**.

Time/NTP

The **Time/NTP** tab enables you to configure system time, date, local timezone, and NTP servers. HP strongly recommends using an NTP server instead of manually configuring the time and date on your system.

Precise time stamping of events is also critical for accurate and reliable log management. The times displayed for Logger operations such as searches, reports, and scheduled jobs are in the Logger's local time zone.

You do not need to configure the time, date, or time zone for a software Logger. Software Loggers use the operating system's settings for the time and time zone.

To set or change the system time, date, or time zone manually:



Caution

If you manually set the date and time settings and are also using an NTP service, the date and time entered manually cannot be more than 16 minutes ahead of or behind the time that the NTP server is providing. If the manually entered time is more than 16 minutes different from the NTP server time, then the NTP service will fail to start.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** in the **System** section.



- 3 In the **Time/NTP** tab, configure these settings.

Setting	Description
Current Time Zone	<p>The time zones appropriate to your system's location. To change this setting, click Change Time Zone...</p> <p>Local times zones follow the Daylight Savings Time (DST) rules for that area. Greenwich Mean Time (GMT) + and - time zones are DST agnostic.</p> <p>For example, the America/Los Angeles time zone varies by an hour compared with GMT when DST goes into and out of effect.</p> <ul style="list-style-type: none"> Pacific Standard Time (PST) = GMT-8 Pacific Daylight Time (PDT) = GMT-7
Current Time	<p>The current date and time at the system's location. To change this setting, click Change Date/Time...</p>

The Time Zone change requires that you reboot the appliance. However, the Current Time change takes effect immediately.

To configure your system as an NTP server or for using an NTP server for your system:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Network** in the **System** section.
- 3 Click the **Time/NTP** tab.
- 4 Under **NTP Servers**, configure these settings.

To add a new NTP server, click the  icon. To remove a server, click the  icon. To change the order in which the NTP servers should be used, select a server and click the up or down arrow until the NTP server is in the desired position.

Setting	Description
Enable as an NTP server	<p>Check this setting if this system should be used as an NTP server.</p>
NTP Servers	<p>Enter the host name of an NTP server. For example, time.nist.gov.</p> <p>HP recommends using at least two NTP servers to ensure precise time on your system. To enter multiple NTP servers, type one server name per line.</p> <p>Once you add servers to this list, you can click the "Click to Test" link to verify if the servers that you added are reachable from your system.</p> <p>Notes:</p> <ul style="list-style-type: none"> An ArcSight system can serve as an NTP server for any other ArcSight system. If System A serves as an NTP server for System B, System B needs to list System A in its NTP Servers list. Use the Test Servers button to verify the status of the servers entered into the NTP Servers box.

- 5 Click **Save**.



You may need to scroll down to view the **Save** button and **Restart NTP Service**.

- 6 Click **Restart NTP Service** to put the changes into effect.

Impact of Daylight Savings Time Change on Logger Operations

Scheduled operations on Logger such as reports, event archives, and file transfers are impacted when system time is adjusted on the Logger at the start and end of the daylight saving time period (DST). The operations scheduled for the hour lost at the start of DST (for example, on March 1, 2012) are not run on the day of time adjustment. Similarly, operations scheduled for the hour gained at the end of the DST (for example, on November 4, 2012) are run at standard time instead of DST time.

Examples:

- A report scheduled to run at 1 a.m. DST on November 4, 2012 will run at 1 a.m. standard time, which is an hour later than the DST time on that day.
- A report scheduled to run at 2 a.m. on November 4, 2012 will run at 2 a.m.; however, due to time adjustment, an hour later than it ran on the previous day (November 3, 2012).
- A report scheduled to run at 2 a.m. on March 11, 2012 will not run.

SMTP

Your system uses the Simple Mail Transfer Protocol (SMTP) setting to send email notifications such as alerts and password reset emails.

To add or change SMTP settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SMTP** in the **System** section and enter these settings.

Setting	Description
Primary SMTP Server	The IP address or hostname of the SMTP server that will process outgoing email.
Backup SMTP Server	The IP address or hostname of the SMTP server that will process outgoing email in case the primary SMTP server is unavailable.
Outgoing Email Address	The email address that will appear in the From: field of outbound email.

- 3 Click **Save**.



Be sure to configure your reports to use the same SMTP settings. For instructions, see ["Report Server Administration" on page 238](#).

License & Update

This page displays license information, the version of the components, and the elapsed time since Logger was last rebooted. From here, you can update Logger and apply a license. To view details of your license, open **Configuration** from the top-level menu bar, and then click **License Information**. For details, see [“License Information” on page 351](#).

Updating the Appliance

To update your Logger:

- 1 Download the update file from the HP Customer Support site (SSO) at <https://support.openview.hp.com> to the computer from which you can connect to Logger.
- 2 Click **System Admin** from the top-level menu bar.
- 3 Click **License & Update** in the **System** section.
- 4 Click **Browse** to locate the file.
- 5 Click **Upload Update**.

An “Update In Progress” page displays the update progress.

- 6 Once the update has completed, the Update Results page displays the update result (success/failure) and whether the update requires a reboot. If the update requires a reboot, the Logger reboots automatically.

Updating the License File

To update a license file:

- 1 Download the update file from the HP Customer Support site (SSO) at <https://support.openview.hp.com> to the computer from which you can connect to the Logger with your browser.

For more information, see [“Acquire a License for the Logger Appliance” on page 34](#).

- 2 From the computer to which you downloaded the update file, log in to the Logger user interface using an account with administrator (upgrade) privileges.
- 3 Click **System Admin** from the top-level menu bar.
- 4 Click **License & Update** in the **System** section.
- 5 Browse to the license file you downloaded earlier, and click **Upload Update**.

An “Update In Progress” page displays the update progress.

Once the update has completed, the Update Results page displays the update result (success/failure). If you are only installing or updating a license, a reboot is not required.

Process Status


The **Process Status** page lists all processes related to your system and enables you to view the details of those processes and start, stop, or restart them.

To view the Process Status page:

- 1 Click **System Admin** from the top-level menu bar.







- Click **Process Status** in the **System** section to display a page similar to the one shown in the following figure.

Process Status




System						
System	Status	Load	CPU Usage	Memory Usage	Data Collected	
n035-h027.qa.arcsight.com	running	[1.89] [1.65] [1.89]	16.8%us 1.9%sy 1.2%wa	68.0% [4076648 kB]	09/15/2010 15:01:23	


NOTE: This Start/Stop buttons are for diagnostic purposes. Please use them with care.

Processes					
  					
Process	Status	Uptime	CPU Usage	Memory Usage	
 apache	running	4h 25m	0.0%	0.0% [3320 kB]	
 aps	running	4h 25m	0.3%	5.0% [303676 kB]	
 connector	running	4h 15m	0.0%	0.0% [608 kB]	

In the **Processes > Process** list on this screen, “processors” refers to forwarders.







- To view the details of a process, click the  icon to the left of the process name, as shown in the following figure.

Process Status



System						
System	Status	Load	CPU Usage	Memory Usage	Data Collected	
n035-h027.qa.arcsight.com	running	[1.50] [1.84] [1.90]	3.9%us 0.5%sy 2.1%wa	68.4% [4097992 kB]	09/15/2010 15:07:54	

NOTE: This Start/Stop buttons are for diagnostic purposes. Please use them with care.

Processes					
  					
Process	Status	Uptime	CPU Usage	Memory Usage	
 apache	running	4h 32m	0.0%	0.0% [3320 kB]	
Children	0				
CPU Percent	0.0%				
CPU Percent Total	0.0%				
Data Collected	09/15/2010 15:08:09				
Memory Kilobytes	3320				
Memory Kilobytes Total	96128				
Memory Percent	0.0%				
Memory Percent Total	1.6%				
Monitoring Status	monitored				
Parent PID	1				
PID	28151				
Status	running				
Uptime	4h 32m				
 aps	running	4h 32m	0.2%	5.1% [311040 kB]	
 connector	running	4h 22m	0.0%	0.0% [608 kB]	

To start, stop, or restart a process, select the process and click **Start**, **Stop**, or **Restart** at the top of the process list.

SNMP

SNMP can be used to push system health events using SNMP traps and poll information from the system. For details on how to poll system health events, see [“Polling System Health Information Using SNMP” on page 366](#). For more information about system health, see [“Monitoring System Health” on page 407](#).

Sending System Health Events as SNMP Traps

You can configure an Alert to send notifications on System Health Events as SNMP traps. Before you do this, you must set up SNMP destinations for them. For more information, see [“Sending Notifications to Syslog and SNMP Destinations” on page 298](#).

Polling System Health Information Using SNMP

You can poll system health information from your system by using SNMP version 2 from any standard network management system.

Refer to your system's Management Information Base (MIB) at `https://<system_name_or_ip>/platform-service/appliance.mib` for the events you can poll from your system. The Object Identifier (OID) **“.1.3.6.1.4.1.11937.3”** identifies the portion of the information tree that contains information relevant for your system.

You can perform the `snmp get`, `getnext`, and `getbulk` operations on your system; `snmp set` operations are not allowed. Additionally, you can perform the `snmp walk` operation to poll for SNMP-related information.

Connecting to Logger to poll system health event information requires a community string. By default, SNMP polling is enabled, and a default, random community string is configured on your system. You can update the default community string to a string that complies with your network security policies. By default, your system listens on port 161 for SNMP requests. You can also update the port number.

Enabling or disabling SNMP polling requires that you reboot your system. However, you do not need to reboot your system when setting a community string or updating the port number.

To enable or disable SNMP polling:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SNMP** in the **System** section.
The SNMP configuration screen displays.

SNMP

The image shows the 'SNMP Configuration' screen. It has a title bar 'SNMP Configuration'. Below it, there are three fields: 'Status' with radio buttons for 'Enabled' (selected) and 'Disabled'; 'Community String' with a text box containing 'Rlk8PKMU44xdL9PaQN9D2hVUQjSNc'; and 'Port' with a text box containing '161'. At the bottom left is a 'Save' button.

- 3 Select **Enabled** to enable and **Disabled** to disable SNMP on your system.
- 4 Click **Save**.

A system message prompts you to reboot your system for the new setting to take effect.



Tip

When moving from a disabled to enabled state, allow at least two minutes before SNMP data is available again.

To change the SNMP community string or port number:

- 1** Click **System Admin** from the top-level menu bar.
- 2** Click **SNMP** in the **System** section.
- 3** Enter the new community string in the Community String field.
- 4** Enter the new port number in the Port field.
- 5** Click **Save**.

Viewing Polled Information

The SNMP information, available through the `snmp get` and `snmp walk` commands, is organized into a sensor table of six columns, as shown in the following table.

Index (1)	Name (2)	Location (3)	Type (4)	Value (5)	Status (6)
An identifier for a row in the sensor table	Name of the sensor	Location of the sensor within the appliance	The type of sensor; for example, fan, power supply, and so on.	Current value of the sensor. Value depends on the appliance model.	Current status of the sensor; for example, OK, Degraded, Rebuilding, Failed, Unavailable
.1.3.6.1.4.1.1 1937.3.1.2.1	Hostname	n/a	Informational	n15-214-128-h92.arst.usa.hp.com	n/a
.1.3.6.1.4.1.1 1937.3.1.2.2	Appliance Model	n/a	Informational	L3200	n/a
.1.3.6.1.4.1.1 1937.3.1.2.4	Cpu0 Usage	cpu0	CPU usage	11%	-
.1.3.6.1.4.1.1 1937.3.1.2.24	Fan1	7.1 (System Board)	Fan	13.72 unspecified	Ok
.1.3.6.1.4.1.1 1937.3.1.2.53	Temp1	64.1 ((Unknown (0x40))	Temperature	26 degrees C	Degraded



The Type and Status of a system are consistent across vendors and various models, however, Name, Value, and Location may vary from vendor to vendor, and across models. Additionally, not all sensors are present in all systems.

HP recommends using any standard MIB browser to view the information obtained via SNMP polling. In a MIB browser, when you perform an SNMP WALK operation, data rows for the first sensor column are returned first, data rows for the second sensor column are returned next, and so on. That is, all the Index sensor entries are returned first, the Name

sensor entries are returned second, the Location sensor entries are returned next, and so on, as shown in the following figure.

Result Table	
Name/OID	
.1.3.6.1.4.1.11937.3.1.1.68	68
.1.3.6.1.4.1.11937.3.1.1.69	69
.1.3.6.1.4.1.11937.3.1.1.70	70
.1.3.6.1.4.1.11937.3.1.1.71	71
.1.3.6.1.4.1.11937.3.1.1.72	72
.1.3.6.1.4.1.11937.3.1.1.73	73
.1.3.6.1.4.1.11937.3.1.2.1	Hostname
.1.3.6.1.4.1.11937.3.1.2.2	Appliance Model
.1.3.6.1.4.1.11937.3.1.2.3	Appliance Version
.1.3.6.1.4.1.11937.3.1.2.4	cpu0 usage
.1.3.6.1.4.1.11937.3.1.2.5	cpu1 usage
.1.3.6.1.4.1.11937.3.1.2.6	cpu2 usage
.1.3.6.1.4.1.11937.3.1.2.7	cpu3 usage
.1.3.6.1.4.1.11937.3.1.2.8	sda bytes read
.1.3.6.1.4.1.11937.3.1.2.9	sda bytes written
.1.3.6.1.4.1.11937.3.1.2.10	FAN MOD 1A RPM
.1.3.6.1.4.1.11937.3.1.2.11	FAN MOD 1B RPM
.1.3.6.1.4.1.11937.3.1.2.12	FAN MOD 2A RPM
.1.3.6.1.4.1.11937.3.1.2.13	FAN MOD 2B RPM
.1.3.6.1.4.1.11937.3.1.2.14	FAN MOD 3A RPM
.1.3.6.1.4.1.11937.3.1.2.15	FAN MOD 3B RPM
.1.3.6.1.4.1.11937.3.1.2.16	FAN MOD 4A RPM
.1.3.6.1.4.1.11937.3.1.2.17	FAN MOD 4B RPM
.1.3.6.1.4.1.11937.3.1.2.18	FAN MOD 5A RPM
.1.3.6.1.4.1.11937.3.1.2.19	FAN MOD 5B RPM
.1.3.6.1.4.1.11937.3.1.2.20	FAN MOD 6A RPM
.1.3.6.1.4.1.11937.3.1.2.21	FAN MOD 6B RPM
.1.3.6.1.4.1.11937.3.1.2.22	Fan Redundancy
.1.3.6.1.4.1.11937.3.1.2.23	RAIDController/Port/p0
.1.3.6.1.4.1.11937.3.1.2.24	RAIDController/Port/p1
.1.3.6.1.4.1.11937.3.1.2.25	Memory Usage
.1.3.6.1.4.1.11937.3.1.2.26	eth0 bytes received
.1.3.6.1.4.1.11937.3.1.2.27	eth1 bytes received
.1.3.6.1.4.1.11937.3.1.2.28	eth0 bytes sent
.1.3.6.1.4.1.11937.3.1.2.29	eth1 bytes sent
.1.3.6.1.4.1.11937.3.1.2.30	eth0 errors
.1.3.6.1.4.1.11937.3.1.2.31	eth1 errors
.1.3.6.1.4.1.11937.3.1.2.32	eth0 link detected
.1.3.6.1.4.1.11937.3.1.2.33	eth1 link detected
.1.3.6.1.4.1.11937.3.1.2.34	RAIDController/Battery/bbu
.1.3.6.1.4.1.11937.3.1.2.35	RAIDController/Configuration
.1.3.6.1.4.1.11937.3.1.2.36	Ambient Temp
.1.3.6.1.4.1.11937.3.1.2.37	Ambient Temp
.1.3.6.1.4.1.11937.3.1.2.38	Ambient Temp

The last two numbers in an OID identify the sensor type and a specific sensor value. In the following OID, .1 identifies the Index sensor type (according to the sensor table described earlier), and .73 identifies the 73rd entry (value) in the index column:

.1.3.6.1.4.1.11937.3.1.1.73

Similarly, in the following figure, .1.3.6.1.4.1.11937.3.1.2.**1**, identifies Hostname, .1.3.6.1.4.1.11937.3.1.2.**2** identifies Appliance Model, .1.3.6.1.4.1.11937.3.1.2.**3** identifies Appliance Version, .1.3.6.1.4.1.11937.3.1.2.**4** identifies CPU0 usage, and so on.

.1.3.6.1.4.1.11937.3.1.1.72	72
.1.3.6.1.4.1.11937.3.1.1.73	73
.1.3.6.1.4.1.11937.3.1.2.1	Hostname
.1.3.6.1.4.1.11937.3.1.2.2	Appliance Model
.1.3.6.1.4.1.11937.3.1.2.3	Appliance Version
.1.3.6.1.4.1.11937.3.1.2.4	cpu0 usage
.1.3.6.1.4.1.11937.3.1.2.5	cpu1 usage
.1.3.6.1.4.1.11937.3.1.2.6	cpu2 usage

To obtain the Value of a sensor, go to the .1.3.6.1.4.1.11937.3.1.5.*n* OIDs. For example, .1.3.6.1.4.1.11937.3.1.5.1 represents the hostname of the system. Similarly, .1.3.6.1.4.1.11937.3.1.5.4 represents 5, the CPU0 usage of the system.

Result Table		
Name/OID	Value	
.1.3.6.1.4.1.11937.3.1.4.60	Voltage	
.1.3.6.1.4.1.11937.3.1.4.61	Voltage	
.1.3.6.1.4.1.11937.3.1.4.62	Voltage	
.1.3.6.1.4.1.11937.3.1.4.63	Voltage	
.1.3.6.1.4.1.11937.3.1.4.64	Voltage	
.1.3.6.1.4.1.11937.3.1.4.65	Voltage	
.1.3.6.1.4.1.11937.3.1.4.66	Voltage	
.1.3.6.1.4.1.11937.3.1.4.67	Voltage	
.1.3.6.1.4.1.11937.3.1.4.68	Voltage	
.1.3.6.1.4.1.11937.3.1.4.69	Voltage	
.1.3.6.1.4.1.11937.3.1.5.1	HP-2106 L2304R2 and www.hp.com	
.1.3.6.1.4.1.11937.3.1.5.2	L3200	
.1.3.6.1.4.1.11937.3.1.5.3	5.3.0.6625.0	
.1.3.6.1.4.1.11937.3.1.5.4	11%	
.1.3.6.1.4.1.11937.3.1.5.5	16%	
.1.3.6.1.4.1.11937.3.1.5.6	13%	
.1.3.6.1.4.1.11937.3.1.5.7	12%	
.1.3.6.1.4.1.11937.3.1.5.8	5151930880	
.1.3.6.1.4.1.11937.3.1.5.9	171493710336	
.1.3.6.1.4.1.11937.3.1.5.10	5640 RPM	
.1.3.6.1.4.1.11937.3.1.5.11	3960 RPM	
.1.3.6.1.4.1.11937.3.1.5.12	5640 RPM	
.1.3.6.1.4.1.11937.3.1.5.13	3840 RPM	
.1.3.6.1.4.1.11937.3.1.5.14	5760 RPM	
.1.3.6.1.4.1.11937.3.1.5.15	3960 RPM	
.1.3.6.1.4.1.11937.3.1.5.16	10560 RPM	
.1.3.6.1.4.1.11937.3.1.5.17	7320 RPM	
.1.3.6.1.4.1.11937.3.1.5.18	3600 RPM	
.1.3.6.1.4.1.11937.3.1.5.19	2760 RPM	
.1.3.6.1.4.1.11937.3.1.5.20	3600 RPM	
.1.3.6.1.4.1.11937.3.1.5.21	2760 RPM	
.1.3.6.1.4.1.11937.3.1.5.22	Fully Redundant	

Follow these steps to view system information in a MIB browser:

- 1 On your appliance:
 - a Access the system MIB from the following URL:
`https://<system_name_or_ip>/platform-service/appliance.mib`
 - b Save it to the system on which you have a MIB browser installed.
- 2 On any standard MIB browser:
 - a Load the MIB. (For example, in the iReasoning MIB Browser, click **File > Load MIBs**, then select the MIB you saved in the previous step, and click **Open**.)
 - b Specify the address and port number of the SNMP agent—your appliance, in this case. (For example, for the iReasoning MIB Browser, enter *IPAddress@port* value in the Address field.)
 - c Configure the community string that is set on your appliance. (For example, in the iReasoning MIB Browser, click **Advanced**, and enter the port number and the Read community string.)
 - d Initiate the snmp WALK operation of the OID from the browser. (For example, in the iReasoning MIB Browser, select **WALK** from the Operations drop-down list.)
 - e Once the SNMP data is returned, interpret it based on the information described earlier in this section.

SSH Access to the Appliance

When you report an issue to customer support that requires them to access your appliance for troubleshooting and diagnostics in situations such as an upgrade failure, unresponsive appliance, and so on, they will direct you to enable SSH access on it.

By default, SSH access (known as Support Login in previous releases) to your appliance is disabled; however, you can select one of these options in the appliance's user interface to enable it:

- Enabled—SSH access is always enabled.
- Enabled, only for 8 hours—SSH access is disabled automatically eight hours after it was enabled.
- Enabled, only during startup/reboot—SSH access is enabled during the time the appliance reboots and is starting up. It is disabled once all processes on the appliance are up and running. This option provides a minimal period of SSH access for situations such as when the appliance does not start successfully after a reboot.



Note

Even if SSH is disabled on your appliance, you can access its console if you have it setup for remote access using the HP ProLiant Integrated Lights-Out (iLO) Advanced remote management card. For more information, see [“Set Up the Logger Appliance for Remote Access”](#) on page 39.

Enabling or Disabling SSH Access

To enable or disable SSH access to your appliance:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSH** in the **System** section.
- 3 Select one of the following options.

SSH

Once you select an option, the user interface displays a message that requires you to confirm the action. Once you confirm it, the change takes effect.

Connecting to Your Appliance Using SSH

Once you have enabled the SSH access, follow these steps to connect to it using SSH:

- 1 Connect to the appliance as “root” using an SSH client.
- 2 When prompted to enter a password, enter a password and press **Enter**.

You are prompted to enter a response to the challenge string displayed on your screen.

- 3 Contact customer support to obtain the challenge response string. Enter it at the "Enter response:" prompt, and press **Enter**.

```
login as: root
root@192.168.36.102's password:
Challenge is SVFNXLCD. Enter response: 1A2B3C4D
```

Logs

Your system can generate audit logs at the application and platform levels. Use the Logs sub-menu to search audit logs.

Audit Logs

Your system's audit logs are available for viewing. Audit logs, as Common Event Format (CEF) audit events, can be sent to ArcSight ESM directly for analysis and correlation. For information about forwarding audit events, see ["Audit Forwarding" on page 371](#).

Audit Logs

Search Audit Logs

Timestamp

12/04/2012

12/04/2012

Description

User

Search

Search Results

User	Description	Timestamp
admin	Successful login	2012/12/04 13:00:21
admin	Authentication Session settings successfully changed.	2012/12/04 11:44:06
admin	Successful login	2012/12/04 11:43:51
admin	Saved Search [SL_Saved_Search_1204_1] has been added	2012/12/04 11:28:06
admin	Saved Search [SL_Saved_Search_1204_2] has been added	2012/12/04 11:28:06
admin	Dashboard [SL_Dashboard_1204_1] has been added	2012/12/04 11:28:06

To view audit logs:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit Logs** in the **Logs** section.
- 3 Select the date and time range for which you want to obtain the log.
- 4 (Optional) To refine the audit log search, specify a string in the Description field and a user name in the User field. When a string is specified, only logs whose Description field contains the string are displayed. Similarly, when a user is specified, only logs whose User field contains the username are displayed.
- 5 Click **Search**.

Audit Forwarding

You can forward audit events to an ArcSight ESM for correlation and analysis. For a list of audit events that you can forward, see [Appendix C, Logger Audit Events, on page 559](#).

To forward audit events to specific ESM destinations:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit Forwarding** in the **Logs** section.
- 3 Select destinations from the **Available Destinations** list and click the right arrow icon (➡) to move the selected destination to the **Selected Destinations** list.

You can select multiple destinations at the same time and move them, or you can move all available destinations by clicking the (➡) icon.

Audit Forwarding

The destinations are ESM destinations that you configure on the ESM Destinations page (**Configuration** (or **Configuration > Settings**) > **Event Input/Event Output > ESM Destinations**).

- 4 Click **Save Settings**.

Storage

Use the Storage sub-menu to add an NFS mount or a CIFS mount, or SAN (if applicable) and to view the status of the hard disk array (RAID) controller and specific system processes.

Remote File Systems

Your system can mount Network File System (NFS) and CIFS (Windows) shares. As a result, it can read log files and event data from UNIX, Linux, Windows remote hosts, and any Network Attached Storage (NAS) solutions based on these operating systems. In addition, you can use the NFS and CIFS mounts for archiving data such as events, exported filters and alerts, and saved searches. Loggers with Storage Area Network (SAN) capability can also interface with a SAN.

Use of NFS for primary storage of Logger events is not recommended. Using a CIFS share for primary storage is not supported. Your system supports only NFS 3.0.






Managing a Remote File System

Make sure the following requirements are met before you mount a share.

File System Type	Requirements
CIFS (Windows)	<ul style="list-style-type: none"> A user account that has access to the shared drive exists on the Windows system. The folder to which you are establishing the mount point is configured for sharing.
NFS	<ul style="list-style-type: none"> Grant your ArcSight system read and write permission on the NFS system. The account used for mounting must use the numeric ids 1500 for uid, or 750 for gid.

To add a Remote File System mount:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Remote File Systems** in the **Storage** section in the left panel. The Remote File Systems form is displayed.

Remote File Systems			
Remote File Systems Configuration			
 Add  Edit  Delete  Test  Refresh			
<input type="checkbox"/> Name ▲	Type	Remote Location	Status
<input type="checkbox"/> nfs_centi	NFS	15.214.157.27:/opt/share	Not In Use
<input type="checkbox"/> sss	CIFS	15.214.157.91:testcifs mount	Not In Use
<input type="checkbox"/> test1	NFS	15.214.130.242:/mnt/volgroup00/qa-testing/cnamineni1	In Use
<input type="checkbox"/> test2	CIFS	15.214.130.194:data	Not In Use
<input type="checkbox"/> testmount	NFS	15.214.130.78:/opt/6642/current/arc sight/logger/logs	In Use

- 3 Click **Add** from the top left side of the page and enter values for the following fields in the resulting form.

Parameter	Description
Select File System Type	Whether you want to mount an NFS or a CIFS share.
NFS Settings	
Name	A meaningful name for the mount point. The name cannot contain spaces. This name is used locally on your system to refer to the mount point, and needs to be specified when configuring archive settings for data that will be stored on the share.
Hostname / IP Address	The name or IP address of the host to which you are creating the mount.
Remote Path (for NFS)	<p>The folder on the remote host that will act as the root of the network file system mount. For example, <code>/public/system_logs</code>.</p> <p>Make sure that only this system can write to the location you specify in this field. If multiple systems (or other systems) mount this location and write to it, data on this location will be corrupted.</p>

Parameter	Description
Mount Options	<p>AutoFS options. For example, <code>ro</code> for read-only from the remote host, <code>rw</code> for read-write, or <code>hard</code> to keep retrying until the remote host responds.</p> <p>Note: Even if you configure <code>rw</code> permission at your mount point, <code>rw</code> permission is not granted to the remote host if the host is configured to allow read-only access.</p>
Description	A meaningful description of the mount point.
CIFS Settings	
Name	A meaningful name for the mount point. The name cannot contain spaces. This name is used locally on your system to refer to the mount point, and needs to be specified when configuring archive settings for data that will be stored on the share.
Location	<p>Enter the share name in one of the following ways:</p> <ul style="list-style-type: none"> Share name in this format: <code><IP Address> or <Hostname>: <share_name></code> For example, <code>198.0.2.160:myshare</code> This folder needs to be configured for sharing. (Typically, to configure a Windows folder for sharing, right click on the folder name > Properties > Sharing.) Caution: when mounting from a Windows Server 2008 in cluster, you must use the Hostname and not the IP address for a successful mount. UNC path For example, <code>//198.0.2.160/myshare</code>
Mount Options	<p>Autofs options. For example, <code>ro</code> for read-only from the remote host, <code>rw</code> for read-write, or <code>hard</code> to keep retrying until the remote host responds.</p> <p>Note: Even if you configure <code>rw</code> permission at your mount point, <code>rw</code> permission is not granted to the remote host if the host is configured to allow read-only access.</p>
Description	A meaningful description of the mount point.
Credentials for CIFS	
Username	<p>The name of the user account with read-write privileges to the Windows share.</p> <p>Make sure the username is prefixed with the domain information. For example, <code>tahoe\arcsight</code>.</p>
Password	The password for the user name specified above.

4 Click **Add**.

All mount points are created under `/opt/mnt`.

To edit a Remote File System mount:**Note**

You cannot edit a mount point if it is in use. The **Edit** link is displayed only if the mount point can be edited.

If you rename a mount point, access to the archives that were made using the original name is lost until you revert the mount point name to the original name.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Remote File Systems** in the **Storage** section in the left panel.
- 3 Select the mount point you want to edit, and click **Edit** from the top left side of the page.
- 4 Change the field values.
- 5 Click **Save**.


To delete a Remote File System mount:**Note**

You cannot delete a mount point that is in use. The **Delete** link is displayed only if the mount point can be deleted.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Remote File Systems** in the **Storage** section in the left panel.
- 3 Select the mount point you want to delete, and click **Delete** from the top left side of the page.

SAN

Some models of the Logger appliance include the ability to connect to a Storage Area Network (SAN). SANs contain Logical Units (LUNs), identified by their World Wide Name. As shown in the following figure, a LUN's Attachment Status can be "available", "attached", or "detached". LUNs in a SAN are in one state at a time. Available actions, such as "attach", vary depending on the state.



The following table summarizes the LUN states and possible actions.

Attachment Status	Actions	Description
available	attach	LUNs detected on a SAN are initially available for attachment.
attached	detach	Attached LUNs can be accessed by Logger. The “ detach ” action is only available if a storage volume has not been configured on the LUN. Once a storage volume has been configured, you cannot “ detach ” the LUN unless you follow the factory reset instructions, described in Appendix H, Restoring Factory Settings , on page 645.
detached	re-attach destroy	When an attached LUN is detached, its data is preserved, but it cannot be accessed by Logger. To make it available again, use the “ re-attach ” action. The “ destroy ” action releases the LUN back to the “ available ” state. When you detach, the only action available immediately is “ re-attach ”. The “ destroy ” state takes a few minutes to appear because it takes a few minutes for the LUN to detach on the system. Destroying a LUN puts it into a state in which a subsequent attach will erase any data stored on the LUN. If a LUN is accidentally destroyed, customer support may be able to recover the data, provided there has been no subsequent attempt to attach the LUN.

Logger can attach to only one LUN at a time for primary storage. You can attach an additional LUN for event archival, configuration backup, and export.

The L7500-SAN has two HBAs. This allows you to use one for multipathing and one for event archival, configuration backup, and export. The L5100-SAN and L7200-SAN have only one HBA. If you set up multipathing on them, the additional LUN is not available. For information about multipathing, see “Creating Multiple Paths to a LUN” on page 458.

Configure SAN settings from the SAN Configuration screen shown in the following figure.

SAN

SAN Configuration							
Refresh							
<input type="checkbox"/> LUN Name	Device	Type	Manufacturer	Mfr. Unique ID	WWN	Size	Status
<input type="checkbox"/> loggerlun	/dev/sde	xfs	3PAR data	352001427	23520002ac001427	536.87 GB	Attached
HBA Information							
HBA 1				HBA 2			
HBA Serial Number	5CF22303NW			HBA Serial Number	5CF22303P8		
HBA Model	HP SN1000E2P 16Gb 2P FC HBA			HBA Model	HP SN1000E2P 16Gb 2P FC HBA		
HBA FW Version	1.0.11.108, sl-4			HBA FW Version	1.0.11.108, sl-4		
HBA Driver Version	Emulex LightPulse Fibre Channel SCSI driver 8.3.5.45.4p			HBA Driver Version	Emulex LightPulse Fibre Channel SCSI driver 8.3.5.45.4p		

To attach a LUN:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN** in the **Storage** section in the left panel.
- 3 Locate and select the LUN in the LUN Name List.

- 4 Click **Attach** from the top left of the SAN Configuration page. If you do not see the **Attach** menu option, no LUNs can be attached to the Logger at this time.



You can attach a LUN only if the LUN is in the "Available" state on Logger.

The LUN's Attachment Status will change to "Attached" when the LUN is ready for use.

To detach a LUN:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN** in the **Storage** section in the left panel.
- 3 In the LUN Name List, locate the LUN to be detached.
- 4 Click **Detach** from the top left of the SAN Configuration page. If you do not see the **Detach** menu option, no LUNs can be detached from the Logger at this time.



You cannot detach a LUN if a storage volume is configured on it.

To re-attach a LUN:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN** in the Storage section in the left panel.
- 3 In the LUN Name List, locate the LUN to be re-attached. The LUN must be in the **Detached** state.
- 4 Click **Re-attach** from the top left of the SAN Configuration page.

If you do not see the Re-attach menu option, no LUNs can be re-attached from the Logger at this time.

To destroy a LUN:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SAN** in the **Storage** section in the left panel.
- 3 In the LUN Name List, locate the LUN to be destroyed. The LUN must be in the 'detached' state.
- 4 Click **Destroy** in the top left corner of the SAN Configuration page.



Destroying a Logical Unit (LUN) that has been detached, puts that LUN into a state in which a subsequent attach will erase any data stored on the LUN. If a LUN is accidentally destroyed, customer support may be able to recover the data, provided there has been no subsequent attempt to attach the LUN.

Restoring a SAN

You can restore a SAN to either the Logger to which it was formerly attached, or to a new Logger (in the case of disaster recovery).

To restore a SAN:

- 1 With Logger powered off, attach the SAN physically.
- 2 Turn on Logger.
- 3 Restore the configuration to Logger. HP recommends backing up the configuration regularly so that a backup file will be available for this purpose. If no backup file is available, skip this step and manually add receivers, forwarders, users, and so on, after the SAN has been restored. For more information, see [“Configuration Backup and Restore” on page 330](#).
- 4 Enable SSH access to your Logger (see [“SSH Access to the Appliance” on page 370](#)).
- 5 Contact customer support at <http://support.openview.hp.com>.
- 6 Customer support will log in remotely, stop all Logger processes, and migrate the internal database to the SAN.
- 7 When customer support has finished, reboot Logger.

Creating Multiple Paths to a LUN

The HBA card on your Logger has two ports. Prior to Logger 5.1, you could connect only one port to the LUN. Starting with 5.1 and later, you can connect both of those ports to the same LUN. Using those ports to create two different paths between the Logger and the LUN (multipathing) reduces the possibility of a single point of failure causing the LUN to become unavailable.



Although any SAN vendor that supports multipathing can work with Logger, ArcSight specifically tests with HP 3PAR SANs.

Logger provides a default multipath configuration as a starting point. However, make sure that you consult your SAN documentation for information specific to your environment.

A multipath user interface (UI) is available by default on Logger models that support SAN. However, you must connect the LUN to both HBA ports and configure multipath configuration in the UI for it to function. Once enabled, **multipath cannot be disabled** on Logger.

Multipath does not enable you to attach an additional LUN to Logger. Only one LUN can be attached at any given time. If multipathing is enabled on your Logger, you cannot use an additional LUN for event archival, configuration backup, and export.

You do not need to enable multipath in order to connect to two different LUNs on different SANs, since there are no duplicate paths. To connect to two different LUNs on the same SAN, or to have two connections to the same LUN, you must configure multipathing. Otherwise, the OS will see duplicate paths to the same LUN, and will be unable to resolve which path to use.

To enable multipath for a new Logger installation, configure multipathing before attaching the LUN. To enable multipath when upgrading from a version prior to Logger 5.1, refer to the release notes for your Logger version.

Enabling Multipath

To enable multipath:

- 1 Ensure that a LUN is **not** attached to the Logger, as described in [“SAN” on page 375](#).

- 2 Click **System Admin** from the top-level menu bar.
- 3 Click **Multipath** in the **Storage** section in the left panel.
- 4 Select a SAN multipathing configuration from the drop-down menu.
- 5 If you chose **Custom**, or if the displayed configuration does not meet your needs, customize the parameters.
- 6 Click **Test** to ensure that the configuration you chose or the changes you made are valid.

If the test fails, make additional changes, or click **Reset** to start over.
- 7 Click **Save**.

When you configure multipath SAN connectivity to the appliance, you must also make sure that the multipathd service is configured to start on boot.

To verify that the multipathd service is configured to start on boot:

- 1 Run `chkconfig --list multipathd`

Make sure '#:on' is shown for your runlevel. The current runlevel can be displayed with the 'runlevel' command.
- 2 If the service is not enabled, do so with:
`chkconfig multipathd on`
- 3 Reboot the appliance or start the multipath daemon with:
`/sbin/service multipathd start`



Be sure to also configure any vendor-specific multipath configuration accordingly in the `/etc/multipath.conf` file.

Restoring Multipath on RMA or Factory Reset Loggers

If you need to restore Logger to its last working state—running version 5.1 or later, with multipath enabled—from 5.0 Patch 3 or earlier, be sure to upgrade your system to Logger 5.1 or later before attaching the LUN.

This could happen in either of the following situations:

- You received a new system that is running Logger 5.0 Patch 3 or earlier after you RMA'd the system to ArcSight.
- You restored the system to its factory default settings, and that reset the Logger version to 5.0 Patch 3 or earlier.

If you will be restoring the configuration of the Logger from a backup, ensure that the Logger is first running the version that matches the backup, perform the restoration, and then complete the upgrade to the desired version. See [Appendix H, Restoring Factory Settings, on page 645](#) for more information.

RAID Controller/Hard Disk SMART Data

You can view information about the RAID controller or hard disk SMART data in the General Controller Information screen. This information is not needed during normal system operations, but it can be helpful for diagnosing specific hardware issues. Due to the redundant nature of RAID storage, a single drive failure will not disable your system.

Instead, performance degrades. Use this report to determine whether a performance issue is caused by a disk failure. Customer support can also use this information to diagnose problems.

To view the General Controller Information screen:

- 1** Click **System Admin** from the top-level menu bar.
- 2** Click **RAID Controller** in the **Storage** section in the left panel.
- 3** The information displayed depends on the hardware model of your system. Click the arrows to open and close the sections.

RAID Controller Configuration

RAID Controller Configuration

▲ **General Controller Information**

```

Bus Interface: PCI
Slot: 0
Serial Number: 5001438011837E30
Cache Serial Number: PBCDH0CRHZT11S
RAID 6 (ADG) Status: Disabled
Controller Status: OK
Hardware Revision: C
Firmware Version: 3.52
Rebuild Priority: Medium
Expand Priority: Medium
Surface Scan Delay: 3 secs
Surface Scan Mode: Idle
Queue Depth: Automatic
Monitor and Performance Delay: 60 min
Elevator Sort: Enabled
Degraded Performance Optimization: Disabled
Inconsistency Repair Policy: Disabled
Wait for Cache Room: Disabled
Surface Analysis Inconsistency Notification: Disabled
Post Prompt Timeout: 15 secs
Cache Board Present: True
Cache Status: OK
Cache Ratio: 25% Read / 75% Write
Drive Write Cache: Disabled
Total Cache Size: 512 MB
Total Cache Memory Available: 400 MB
No-Battery Write Cache: Disabled
Cache Backup Power Source: Capacitors
Battery/Capacitor Count: 1
Battery/Capacitor Status: OK
SATA NCQ Supported: True
          
```

▼ **Logical Drive #1**

▼ **Disk #0**

▼ **Disk #1**

▼ **Disk #2**

Security

Security settings enable you to configure SSL server certificates, enable and disable FIPS (Federal Information Processing Standards) mode on your system, and configure SSL client authentication for client certificate and Common Access Card (CAC) support.



For steps on how to create a user DN, see [“Users” on page 402](#), and refer to the section “Use Client DN” in the parameters table.

SSL Server Certificate

Your system uses Secure Sockets Layer (SSL) technology to communicate securely over an encrypted channel with its clients, such as SmartConnectors, when using the SmartMessaging technology and other ArcSight systems. Your system ships with a self-signed certificate so that an SSL session can be established the first time you use the appliance. For more information on this option, see [“Generating a Self-Signed Certificate” on page 381](#).

Although a self-signed certificate is provided for your use, HP strongly recommends using a certificate authority (CA) signed certificate. Even if FIPS is not enabled on your system, it must use a **CA-signed certificate** if it is a destination of a FIPS-enabled container (on a Connector Appliance) or a software-based SmartConnector. Additionally, ensure that the root certificate of the CA that signed your system's certificate is trusted on the container or the SmartConnector. If the CA's root certificate is not trusted, load it on the container by following instructions in the “Managing Certificates on a Container” section in the Connector Appliance Administrator's Guide.

To facilitate obtaining a CA-signed certificate, your system can generate a Certificate Signing Request. Once a signed certificate file is available from the CA, it can be uploaded to your system for use in a subsequent authentication. For detailed instructions, see [“Generating a Certificate Signing Request \(CSR\)” on page 383](#).

Your system generates an audit event when the installed SSL certificate is going to expire in less than 30 days or has already expired. The event with Device Event Class ID “platform:407” is generated periodically until you replace the certificate with one that is not due to expire within 30 days.

Generating a Self-Signed Certificate

Your appliance ships with a self-signed certificate so that an SSL session can be established the first time you connect. This type of certificate does not require signing from another entity and can be used immediately.

To generate a self-signed certificate:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** from the **Security** section in the left panel to display the **Generate Certificate/Certificate Signing Request** page.

- 3 Click the **Generate Certificate** tab.

SSL Settings

Generate Certificate | **Import Certificate**

Generate Certificate/Certificate Signing Request

Enter Certificate Settings

Country (2-letter code)	US
State/Province	California
City/Locality	Sunnyvale
Organization Name	Hewlett-Packard
Organizational Unit	Support Team
Hostname	scm012.hewlett-packard.hp.com
Email Address	arst-support@hp.com
Private Key Length	1024

Generate CSR **Generate Certificate** **View Certificate**

- 4 From the **Enter Certificate Settings** field, enter new values for the following fields:

Parameter	Description
Country	A two-letter country code, such as 'US' for the United States.
State/Province	State or province name, such as 'California.'
City/Locality	City name, such as 'Sunnyvale'.
Organization Name	Company name, governmental entity, or similar overall organization.
Organizational Unit	Division or department within the organization.
Hostname	<p>The host name or IP address of this system.</p> <p>When specifying the host name, make sure that this name matches the name registered in the Domain Name Service (DNS) server for the system. Additionally, this name must be identical to the host name specified in "NICs" on page 359.</p> <p>Note: If the host name or IP address of this system changes in the future, you must generate a new self-signed certificate or CSR. Once a new certificate is obtained, you must upload it to ensure that the connectors (in FIPS mode) which communicate with the system are able to validate the host name.</p>
Email Address	The email address of the administrator or contact person for this CSR.
Private Key Length	Select the length (in bits) of the private key: 1024 , 2048 , 4096 , or 8192 .

Use the first two buttons to generate a CSR or a self-signed certificate. The **View Certificate** button is only used to view the resulting certificate.

Button	Description
Generate CSR	Click to generate a Certificate Signing Request (CSR).
Generate Certificate	Click to generate a self-signed certificate.
View Certificate	Click to view the generated certificate.

- 5 Click the **Generate Certificate** button to generate the self-signed certificate.



The Apache server restarts while generating the certificate. You may get an error communicating to the web server while this is happening. This is expected behavior, and communication is automatically restored once Apache is back up.

- 6 Click **Ok** after the confirmation message appears.
- 7 Click the **View Certificate** button to view the PEM encoded self-signed certificate.

Generating a Certificate Signing Request (CSR)

The first step in obtaining a CA-signed certificate is to generate a Certificate Signing Request (CSR). The CSR must be generated on the system for which you are requesting a certificate. That is, you cannot generate a CSR for System A on System B or use a third-party utility for generation.

The resulting CSR must be sent to a CA, such as VeriSign, which responds with a signed certificate file.

To generate a certificate signing request:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** from the **Security** section in the left panel to display the **Generate Certificate/Certificate Signing Request** page.
- 3 Click the **Generate Certificate** tab.

SSL Settings

Generate Certificate Import Certificate

Generate Certificate/ Certificate Signing Request

Enter Certificate Settings

Country (2-letter code) US

State/Province California

City/Locality Sunnyvale

Organization Name Hewlett-Packard

Organizational Unit Support Team

Hostname scm001c.hpe.com

Email Address arst-support@hp.com

Private Key Length 1024

Generate CSR Generate Certificate View Certificate

- 4 From the **Enter Certificate Settings** field, enter new values for the following fields:

Parameter	Description
Country	A two-letter country code, such as 'US' for the United States.
State / Province	State or province name, such as 'California.'
City / Locality	City name, such as 'Sunnyvale'.
Organization Name	Company name, governmental entity, or similar overall organization.
Organizational Unit	Division or department within the organization.
Hostname	<p>The host name or IP address of this system.</p> <p>When specifying the host name, make sure that this name matches the name registered in the Domain Name Service (DNS) server for the system. Additionally, this name must be identical to the host name specified in "NICs" on page 359.</p> <p>Note: If the host name or IP address of this system changes in the future, you must generate a new self-signed certificate or CSR. Once a new certificate is obtained, you must upload it to ensure that the connectors (in FIPS mode) which communicate with the system are able to validate the host name.</p>
Email Address	The email address of the administrator or contact person for this CSR.
Private Key Length	Select the length (in bits) of the private key: 1024 , 2048 , 4096 , or 8192 .

Use the first two buttons to generate a CSR or a self-signed certificate. The **View Certificate** button is only used to view the resulting certificate.

Button	Description
Generate CSR	Click to generate a Certificate Signing Request (CSR).
Generate Certificate	Click to generate a self-signed certificate.
View Certificate	Click to view the generated certificate.

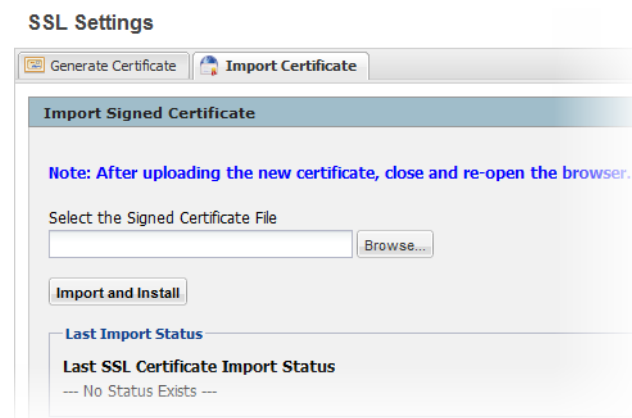
- 5 Choose **Generate CSR** to generate a certificate signing request.
- 6 If the CSR was successfully generated, a pop-up window appears, allowing you to either download the CSR file or to cut-and-paste its content.

To do so, copy all the lines from -----BEGIN CERTIFICATE REQUEST----- to -----END CERTIFICATE REQUEST-----.
- 7 Send the CSR file to your certificate authority to obtain the CA-signed certificate.
- 8 Once the CA-signed certificate file is obtained, continue on to [Importing a Certificate](#) below.

Importing a Certificate

If you have obtained a certificate from your certificate authority (CA), follow the steps below to import it onto your system.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** under the **Security** section in the left panel.
- 3 Select the **Import Certificate** tab.



- 4 Click the **Browse** button to locate the signed certificate file on your local file system.



The imported certificate must be in **Privacy Enhanced Mail (PEM)** format.

Note

- 5 Click **Import and Install** to import the specified certificate.

- 6 If using **HTTPS** and depending on your browser, you may need to close and restart the browser for the new certificate to take effect. If you are unsure of your browser's requirements, close and restart it.

SSL Client Authentication

Your system supports client authentication using SSL certificates. SSL client authentication is a form of two-factor authentication that can be used *as an alternate* or *in addition to* local password authentication. As a result, your system can be configured for SmartCards, such as Common Access Card (CAC) based authentication. CAC is a standard identification card for active duty members of the Uniformed Services, Selected Reserve, DOD civilian employees, and eligible contractor personnel.



Note

CAC is a form of client certificate authentication. Information on client certificate authentication applies to CAC.

Your system also supports LDAPS authentication. The SSL certificate for the LDAPS server must be uploaded into the trusted store. After uploading the SSL certificate, the aps process must be restarted (**System Admin > Process Status > aps > Restart**).

Configuring Logger to Support SSL Client Authentication

Perform the following steps to configure Logger to support SSL client authentication.

On the Logger:

- 1 If the Logger uses the default signed certificate it shipped with from ArcSight, replace it with a *FIPS-compliant*, signed SSL server certificate. Follow instructions at ["Uploading Trusted Certificates" on page 387](#) to load the certificate.



Caution

All SSL client certificates used for authentication must be FIPS compliant (that is, hashed with FIPS-compliant algorithms) even if FIPS is not enabled on your Logger.

- 2 Enable client certificate authentication, as described in ["Client Certificate Authentication" on page 397](#).
- 3 If the client certificates are **CA signed**, upload the root certificate of the authority who signed the certificates that will be used for authenticating clients, as described in ["Uploading Trusted Certificates" on page 387](#).

If the client certificates used to authenticate with Logger are signed by different CAs, make sure you upload root certificates of **all** CAs.

If the client certificates are **self-signed**, upload the public portion of the client certificate.

- 4 Configure a user name for each user who will be connecting to the Logger using a client certificate, as described in ["User Management" on page 402](#).
- 5 (Optional) Upload a certificate revocation list (CRL), as described in ["Uploading a Certificate Revocation List" on page 387](#).
- 6 (Optional) If this Logger is configured to use **only** SSL Client Authentication, make sure this Logger's Authorization ID and Code are appropriately configured on other Loggers that with it. For more information, see ["Peer Loggers" on page 326](#).

On the Client (Web browser):

Configure your browser to provide the SSL client certificate when accessing Logger.
(Upload the private key in PKCS 12 format in your browser.)

Uploading Trusted Certificates

A trusted certificate is used to authenticate users that log in to your system. Uploading a trusted certificate is required if you are using LDAPS authentication. The trusted certificate is used to authenticate the remote LDAPS server. The certificate needs to be in Privacy Enhanced Mail (PEM) format.

To upload a trusted certificate:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Client Authentication** in the **Security** section in the left panel.
- 3 On the **Trusted Certificates** tab, click **Browse** to find the trusted certificate on your local file system.
- 4 Click **Upload**.

The trusted certificate is uploaded and listed in the “Certificates in Repository” list on the same page where you uploaded it.

To view details about a trusted certificate, click the link displayed in the Certificate Name column.

To delete a trusted certificate, select the certificate and click **Delete**.

Uploading a Certificate Revocation List

A certificate revocation list (CRL) is a computer-generated record that identifies certificates that have been revoked or suspended before their expiration dates. To support CAC, you need to upload a CRL file to your ArcSight system. The CRL file needs to be in PEM format.

To upload a CRL file:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Client Authentication** in the **Security** section in the left panel.

- 3 In the **Certificate Revocation List** tab, click **Browse** to find the CRL file on your local file system.
- 4 Click **Upload**.

The CRL is uploaded and listed in the Certificate Revocation List.

To view details about a CRL, click the link displayed in the Issuer Name column.

To delete a CRL file, select it and click the **Delete** button.

Enabling Client Certificate Authentication

To enable client certificate authentication, see [“Client Certificate Authentication” on page 397](#).

FIPS 140-2

Your system supports the Federal Information Processing Standard 140-2 (FIPS 140-2). FIPS 140-2 is a standard published by the National Institute of Standards and Technology (NIST) and is used to accredit cryptographic modules in software components. The US Federal government requires that all IT products dealing with Sensitive, but Unclassified (SBU) information meet these standards.

If your system needs to be FIPS 140-2 compliant, you can enable FIPS. Once you do so, the system uses the cryptographic algorithms defined by the NIST for FIPS 140-2 for all encrypted communication between its internal and external components.



Note

To be fully FIPS 140-2 compliant, all components of your Logger deployment need to be in FIPS 140-2 mode. For example, if you enable FIPS 140-2 on your Logger but the SmartConnectors that send events to it are not running in FIPS 140-2 mode, your deployment is not fully FIPS 140-2 compliant.

In a typical deployment, your Logger will communicate with the following components. To be fully FIPS compliant, all of these components should be FIPS enabled:

- SmartConnectors that send events to it

FIPS mode is supported on SmartConnectors running version 4.7.5.5372 and later. Follow instructions in [“Installing or Updating a SmartConnector to be FIPS compliant” on page 390](#) to ensure that your connector is FIPS compliant.

- **Logger forwarders, such as ArcSight Managers to which Logger forwards events and alerts**

The system to which your FIPS-compliant Logger forwards events should be FIPS compliant as well. Additionally, you need to import that system's SSL server certificate on the Logger so that Logger can communicate with it.

If you forward events and alerts to an ArcSight Manager, it needs to run ESM 4.0 SP2 or later to enable FIPS 140-2 on it. For more information, see the ArcSight ESM Installation and Configuration Guide for the ESM version you are running. Additionally, follow instructions in [“ESM Destinations” on page 291](#) to complete configuration of this setup.

- **Loggers**

Loggers running 45.0 or later automatically use FIPS 140-2 compliant algorithms. Therefore, no action is required on such Loggers, except enabling FIPS as described in this section.

- **Connector Appliance**

If your Logger platform includes an integrated Connector Appliance, both products operate in FIPS mode when you enable FIPS on the Logger. However, you might need to do additional configuration on the Connector Appliance components for FIPS-mode operation. See the Connector Appliance Administrator's Guide for more information.

A Logger must use a CA-signed certificate if it is a destination of a FIPS-enabled container (on a Connector Appliance) or a software-based SmartConnector. Additionally, ensure that the root certificate of the CA that signed Logger's certificate is trusted on the container or the SmartConnector. If the CA's root certificate is not trusted, load it on the container by following instructions in the “Managing Certificates on a Container” section in the Connector Appliance Administrator's Guide.

You can enable or disable FIPS mode on Logger to suit your needs; however, you will need to reboot the appliance before the new mode will be effective. If your Logger platform has an integrated Connector Appliance, make sure you have read the FIPS 140-2 information specific to the Connector Appliance in the Connector Appliance Administrator's Guide before disabling FIPS.

Things to be Aware of When Enabling FIPS Mode on Logger:

- Your Logger must be set up with a CA-signed SSL certificate. For more information, see [“SSL Server Certificate” on page 381](#).
- A Logger, even when in non-FIPS mode, must use a CA-signed certificate if it is a destination of a FIPS-enabled container (on a Connector Appliance) or a software-based SmartConnector. Additionally, ensure that the root certificate of the CA that signed Logger's certificate is trusted on the container or the SmartConnector. If the CA's root certificate is not trusted on the container, load it by following instructions in the “Managing Certificates on a Container” section in the Connector Appliance Administrator's Guide. If the CA's root certificate is not trusted on the SmartConnector, follow instructions in [“Installing or Updating a SmartConnector to be FIPS compliant” on page 390](#).
- Once FIPS is enabled on your Logger, the SmartMessage receiver (if configured) stops receiving events from non-FIPS connectors if those connectors are not running version 4.7.5.5372 or later. Make sure you have the correct connectors.

To enable or disable FIPS mode:

Make sure you are familiar with the configuration requirements on your Logger as described in “Things to be Aware of When Enabling FIPS Mode on Logger:” on page 389.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **FIPS 140-2** in the Security section in the left panel.
- 3 Click **Enable** or **Disable** for the Select FIPS Mode option.
- 4 Click **Save**.
- 5 Reboot your Logger.

The FIPS Status Table shows which processes and components of the Logger are FIPS enabled.

Installing or Updating a SmartConnector to be FIPS compliant

The information in this section is same as that in the ArcSight Installing FIPS-Compliant SmartConnectors document except that the information in that document is generally applicable, while information in this section is in the context of Logger.

FIPS mode is supported on SmartConnectors running version 4.7.5.5372 or later.

If you are...	Then...
Installing a new SmartConnector to send events to a Logger in FIPS-compliant mode	<ol style="list-style-type: none"> 1 Download a FIPS-supported SmartConnector version from the HP Customer Support site (SSO). 2 Go to Step 1 on page 390.
Updating a SmartConnector to be FIPS compliant and the SmartConnector is not running version 4.7.5.5372 or later.	<ol style="list-style-type: none"> 1 Upgrade the SmartConnector to a FIPS-supported version. Follow instructions in the SmartConnector User's Guide to upgrade the SmartConnector. 2 Perform only Step 2a on page 391.
Updating a SmartConnector to be FIPS compliant and the SmartConnector is running version 4.7.5.5372 or later.	Perform only Step 2a on page 391 .

To make a SmartConnector FIPS compliant:

- 1 Follow device configuration steps provided in the SmartConnector's configuration guide (available from the HP Customer Support site (SSO) at <http://support.openview.hp.com>), then follow the installation procedure through installation of the core connector software (SmartConnector Installation step 2).

At Step 3 of the Connector setup, click **Cancel** to exit the setup. You must then configure the NSS DB, which is necessary for installing the connector in FIPS-compliant mode.

Once the NSS DB is configured, continue with [Step 2](#), below.

- 2 To enable FIPS Mode on the SmartConnector:

- a** Create an `agent.properties` file at the following location if it does not exist already:

```
$ARCSIGHT_HOME\current\user\agent
```

- b** Enter the following property, then save and close the file.

```
fips.enabled=true
```

3 Import Logger's Certificate on the SmartConnector:

- a** In a DOS prompt window on your SmartConnector machine, from `$ARCSIGHT_HOME\current\bin`, enter the following command to turn off FIPS mode.

```
arcsight runmodutil -fips false -dbdir  
user/agent/nssdb.client
```

- b** Export the Logger certificate file and import it to the SmartConnector's NSS DB as follows:

- i** Export Logger's certificate file from the browser you use to connect to it. Refer to your browser's Help for instructions. For example, to export a Logger's certificate file on Firefox, click **Tools > Options > Advanced > Encryption > View Certificates > Servers > *Select your Logger appliance* > Export**. Save the certificate file with a `.crt` or `.cer` extension.

- ii** Copy the certificate file you exported in the previous step (in this example, **loggercert.crt**) to the `$ARCSIGHT_HOME\current\bin` directory on the SmartConnector. From `$ARCSIGHT_HOME\current\bin`, enter the following:

```
arcsight runcertutil -A -n mykey -t "CT,C,C" -d  
user/agent/nssdb.client -i bin/loggercert.crt
```

- c** Enter the following command to re-enable FIPS mode that you turned off in Step 1:

```
arcsight runmodutil -fips true -dbdir  
user/agent/nssdb.client
```

- d** Ensure that the SmartConnector can resolve the name specified in the CN value of the Logger certificate's *Subject* field. If the name is not resolvable, add it to SmartConnector system's Hosts file.

- e** If you are installing a new SmartConnector, continue to the next step.

If you are updating your SmartConnector to be FIPS compliant, ensure that the connector's Logger destination host name is same as the CN value in the certificate's *Subject* field, and **exit this procedure**.

- 4** To return to the SmartConnector configuration wizard, enter the following from `$ARCSIGHT_HOME\current\bin`:

```
arcsight connectorsetup
```

- 5** When prompted whether you want to start in Wizard Mode, click **Yes**.

The Destination selection window is again displayed; return to your SmartConnector Configuration Guide, **SmartConnector Installation step 4** to continue the connector configuration.



Note

When configuring the connector, ensure that the connector's Logger destination host name is same as the CN value in the certificate's **Subject:** field.

For the remainder of the configuration process, see the Configuration Guide for the SmartConnector you selected to install. The specific configuration guide provides information about how to configure the device for event collection, specific installation parameters required during the configuration process, and a table of vendor-specific field mappings to ArcSight events.

Users/Groups

Use the **Users/Groups** sub-menu to configure users and user groups, and to set authentication options.

Authentication

Authentication Settings enable you to specify the settings and policies for user login sessions, password rules and lockouts, and external authentication options.

Sessions

The **Session** tab lets you specify the maximum number of simultaneous sessions for a single user account, and the length of time after which a user session is automatically logged out or a user account disabled. By default, a single user account can have up to 15 simultaneous active sessions, and a user account is logged out after 15 minutes of inactivity.

Authentication Settings

Sessions Local Password External Authentication

Session Settings

Max Simultaneous Logins/User: 15

Logout Inactive Session After: 0 hours 15 minutes 0 seconds

☒ Disable Inactive Account After: 0 days

Save

To change session settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 On the **Sessions** tab, update the parameters described in the following table.

Parameters	Description
Max Simultaneous Logins/User	The maximum number of simultaneous sessions allowed for a single user account. The default is 15 sessions .

Parameters	Description
Logout Inactive Session After	The length of time, in minutes, after which an inactive session is automatically ended. The default is 15 minutes . This value does not apply to the user interface pages accessed through the Monitor menu. If a user is on any of the Monitor menu pages and the session has been inactive for the specified number of minutes, the user's session remains active.
Disable Inactive Account After	The number of days after which an inactive user account is disabled. The default is 0 , meaning the account is never disabled.

- 4 Click **Save** to make the changes, or click another tab to cancel.

Local Password

The **Local Password** tab enables you to set password policies, such as the minimum and maximum number of characters and other password requirements.

The screenshot shows the 'Authentication Settings' window with the 'Local Password' tab selected. The 'Local Password Settings' section contains three main areas: 'Lockout Account', 'Password Expiration', and 'Password Strength Rules'. The 'Lockout Account' section has checkboxes for 'Enable Account Lockout' (checked), 'Remember Failed Attempts For' (0 hours, 1 minutes, 0 seconds), and 'Lockout Account For' (0 hours, 15 minutes). The 'Password Expiration' section has checkboxes for 'Enable Password Expiration' (checked), 'Password Expires in' (90 days), and 'Notify User' (5 Days Before Expiration), with a link to 'Users Exempted From Password Expiration Policy'. The 'Password Strength Rules' section has a checkbox for 'Enforce Password Strength' (checked), 'Minimum Length' (10 characters), 'Maximum Length' (20 characters), and a 'Password Character Rules' section with checkboxes for 'Numeric [0-9]' (2), 'Uppercase [A-Z]' (0), 'Special [!\$%^*...]' (2), 'Lowercase [a-z]' (0), and 'Password Must be At Least' (2 Characters Different From Old Password). At the bottom, there is a checkbox for 'Include "Forgot Password" link on Login Screen' and a 'Save' button.

To change the password settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **Local Password** tab.

Use the parameters described in the following table to customize your password settings.

Table 7-1 Authentication Settings, Local Password tab

Parameter	Description
Lockout Account (policy)	
Enable Account Lockout	Select the checkbox to enable user accounts to be locked out as defined by the following settings. By default, the policy is disabled .
Lockout Account After	Number of failed login attempts after which a user account is locked out. The default is 3 .
Remember Failed Attempts For	The length of time, in minutes, for which a failed login attempt is remembered. The default is 1 .
Lockout Account For	The length of time, in minutes, for which a locked out account cannot be unlocked. The default is 15 .
Password Expiration (policy)	
Enable Password Expiration	Select the checkbox to enable user passwords to expire as defined by the following settings. By default, the policy is disabled .
Password Expires in	Number of days after which the password expires. The default is 90 .
Notify User	Number of days before expiration to notify the user. Select this option to allow users to update their password before expiration. The default is 5 .
Users Exempted From Password Expiration Policy	Click the link to set the number of users whose password should never expire. For information on how to use this feature, see "Users Exempted From Password Expiration" on page 395.
Password Strength Rules (policy)	
Enforce Password Strength	Select the checkbox to enforce password policy as defined by the following settings. By default, the policy is disabled .
Minimum Length	Minimum number of characters that a password must contain. The default is 10 .
Maximum Length	Maximum number of characters that a password can contain. The default is 20 .
Password Character Rules	
Password character rules define additional character requirements to ensure password strength.	
Numeric	Minimum number of numeric characters (0-9) in a password. The default is 2 .
Uppercase	Minimum number of uppercase characters (A-Z) in a password. The default is 0 .
Special	Minimum number of non-digit and non-letter characters that are required in a password. The default is 2 .
Lowercase	Minimum number of lowercase characters (a-z) in a password. The default is 0 .

Table 7-1 Authentication Settings, Local Password tab (Continued)

Parameter	Description
Password Must be At Least N Characters Different From Old Password	Minimum number of characters by which the new password must differ by from the previous one. The default is 2 .
Include "Forgot Password" link on Login Screen	<p>Select the checkbox to enable users to reset their local password via a "Forgot Password" link on the login page. By default, the option is disabled.</p> <p>An SMTP server must be configured on the system, and the username must have a correct email address for this feature to work successfully.</p> <p>If an SMTP server is not set, you will not be able to reset the password because the email containing the temporary password cannot be sent.</p> <p>An email address must be specified in the user settings for the user name. The temporary password is sent to that email address. If no email address is specified or if the email address is incorrect, the user will not receive the email.</p> <p>For information on how to use this feature, see "Forgot Password" on page 396.</p>


- 4 Click **Save** to save the changes, or click another tab to cancel.


Users Exempted From Password Expiration


Even though you have set a password expiration policy for most users, you may want to have a user whose password does not expire automatically.

To exempt a user from the password expiration policy:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **Local Password** tab, and then click **Users Exempted From Password Expiration Policy**.
- 4 The **Exempt Users From Password Expiration** page appears.



- 5 Select users from the **Non-exempted Users** list and click the right arrow icon  to move the selected users to the **Exempted Users** list. Do the reverse to remove users from the list of exempted users.

You can select multiple users at the same time and move them over. Or you can move all users by clicking the  icon.

- 6 Click **Save** to save the policy or **Cancel** to exit.

Forgot Password

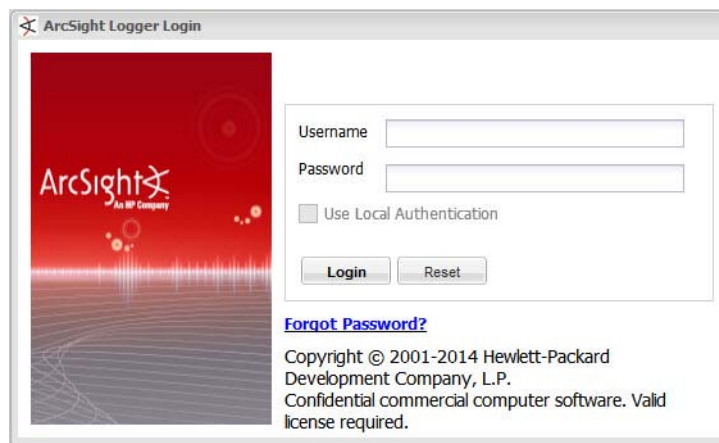
This feature is available only if the “Include “Forgot Password” link on Login Screen” setting on the Authentication Settings page (**System Admin > Authentication > Local Password**) is set to **Yes**. By default, this setting is set to **No**. An SMTP server must be configured in order to use this feature. For more details on how to enable it, see [“Local Password” on page 393](#).

If you forget your system password, use this feature to receive an email that provides a temporary password.

The temporary password is valid until the time specified in the email. If you do not log in within the specified time, only an administrator can reset the password to generate another temporary password.

To reset your password:

- 1 Click the **Forgot Password** link on the Login screen, as shown in the following figure.



- 2 Enter a user name on the Reset Password screen.
- 3 Click **Reset Password**.

An automated email with a temporary password is sent to the email address specified for that user.

External Authentication

Besides providing a local password authentication method, your system supports Client Certificate/CAC, LDAP, and RADIUS authentication. It is not possible to enable all authentication methods simultaneously.



Note

CAC is a form of client certificate authentication. Information on client certificate authentication applies to CAC.

From the **External Authentication** tab, use the drop-down menu to choose one of the following authentication methods:

- [Local Password](#)
- [Client Certificate Authentication](#)
- [Client Certificate and Local Password Authentication](#)
- [LDAP/AD and LDAPS Authentication](#)
- [RADIUS Authentication](#)

Local Password

This option is the default method and implements the local password policies set in the **Local Password** tab. Leave this as the default, or click **Save** if changing from another option.

Client Certificate Authentication

This authentication method requires that users authenticate using a client certificate. For each client certificate, a user account with a Distinguished Name (DN) matching the one in the client certificate must exist on your system.



All SSL client certificates used for authentication must be FIPS compliant (hashed with FIPS-compliant algorithms) even if FIPS is not enabled on your system.

To configure client certificate authentication:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **External Authentication** tab.
- 4 From the drop-down menu, choose **Client Certificate**.
- 5 **Allow Local Password Fallback** provides two options:
 - ◆ **Allow Local Password Fallback for Default Admin Only**
Select this option to allow the default admin user to log in using only a username and password if the client certificate is not available or invalid. This privilege is restricted to the default admin user only—other users must have a valid client certificate to gain access to the system. This option is enabled by default.
 - ◆ **Allow Local Password Fallback for All Users**
Select this option to allow all users to log in using their local user name and password if their client certificate is invalid or unavailable.

For more information, see [“Local Password Fallback” on page 401](#).
- 6 Click **Save**.

Client Certificate and Local Password Authentication

This authentication method requires that users authenticate using an SSL client certificate and a valid local password. *Local Password* refers to the password associated with the user credentials created in **User Management** in the **Users/Groups** section. See [“User Management” on page 402](#) for details.

A user account on your system must be defined with a Distinguished Name (DN) that matches the one in the client certificate.

For instructions on how to create a user DN, see [“Users” on page 402](#) and refer to the section called “Use Client DN” in the parameters table.



All SSL client certificates used for authentication must be FIPS compliant (hashed with FIPS-compliant algorithms) even if FIPS is not enabled on your system.

To configure client certificate and password authentication:

- 1 Click **System Admin** from the top-level menu bar.
 - 2 Click **Authentication** in the **Users/Groups** section.
 - 3 Choose the **External Authentication** tab.
 - 4 From the drop-down menu, choose **Client Certificate AND Local Password**.
 - 5 **Allow Local Password Fallback** provides two options:
 - ◆ **Allow Local Password Fallback for Default Admin Only**
This option, always enabled, allows the default admin user to log in using only a username and password.
 - ◆ **Allow Local Password Fallback for All Users**
This option is always disabled. You cannot enable it when using the **Client Certificate AND Local Password** authentication method.
- For more information, see [“Local Password Fallback” on page 401](#).
- 6 Click **Save**.

LDAP/AD and LDAPS Authentication

This authentication method authenticates users against an LDAP server. Even when LDAP is enabled, each user account must exist locally on your system. Although the user name specified locally can be different from the one specified on the LDAP server, the Distinguished Name (DN) specified for each user account must match the one in the LDAP server.



For steps on how to create a user DN, see [“Users” on page 402](#), and the parameter [“Use Client DN” on page 403](#).

To set up LDAP authentication:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **External Authentication** tab.
- 4 From the drop-down menu, choose **LDAP**.
- 5 **Allow Local Password Fallback** provides two options:

◆ **Allow Local Password Fallback for Default Admin Only**

Select this option to allow the default admin user to log in using only a username and password if LDAP authentication fails. This privilege is restricted to the default admin user only—all others must be authenticated by LDAP. This option is enabled by default.

◆ **Allow Local Password Fallback for All Users**

Select this option to allow all users to log in using their local user name and password if LDAP authentication fails.

For more information, see [“Local Password Fallback” on page 401](#).

LDAP Server has the following parameters:

Parameter	Description
Server Hostname[:port] (optional)	<p>(Optional) Enter the host name or IP address and port of the LDAP server in the following format:</p> <pre>ldap://<hostname or IP address>:<port></pre> <pre>ldaps://<hostname or IP address>:<port></pre> <p>Additional steps are required for the use of LDAPS. See Using the LDAP over SSL (LDAPS) Protocol below.</p>
Backup Server Hostname[:Port] (optional)	<p>(Optional) Enter the backup LDAP server to use if the primary server does not respond. If the server returns an authentication failure (bad password, unknown username, etc), then the backup server is not tried. The backup server is tried only when the primary server has a communication failure.</p> <p>Use the same format as the primary server to specify the host name and port.</p>
Request Timeout	The length of time, in seconds, to wait for a response from the LDAP server. The default is 10 .

6 When finished, click **Save**.

Using the LDAP over SSL (LDAPS) Protocol

When choosing the LDAPS protocol to authenticate users, make sure the following conditions are true:

- The SSL certificate for the LDAPS server has been uploaded into the trusted store.
- The external authentication method is set to “LDAP”.
- The URL for the LDAPS server(s) starts with “ldaps://”.

After uploading the SSL certificate, the **aps** process must be restarted (**System Admin > Process Status > aps Restart**).



If the aps process is not restarted, attempts to authenticate via LDAPS will fail.

RADIUS Authentication

This authentication method allows users to authenticate against a RADIUS server. Even when RADIUS authentication is enabled, each user account must exist locally on your system. The username must match the one in the RADIUS server, although the password can be different. A user must present a valid username and (RADIUS) password to be successfully authenticated.

To configure RADIUS authentication settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **External Authentication** tab.
- 4 From the drop-down menu, choose **RADIUS**.
- 5 **Allow Local Password Fallback** provides two options:
 - ◆ **Allow Local Password Fallback for Default Admin Only**
Select this option to allow the default admin user to log in using only a username and password if RADIUS authentication fails. This privilege is restricted to the admin user only—all others must be authenticated by RADIUS. This option is enabled by default.
 - ◆ **Allow Local Password Fallback for All Users**
Select this option to allow all users to log in using their local user name and password, if RADIUS authentication fails. For more information, see [“Local Password Fallback” on page 401](#).
- 6 **Update the RADIUS Server** parameters as necessary:

Parameter	Description
Server Hostname[:port]	Enter the host name and port of the RADIUS server.
Backup Server hostname[:port] (optional)	(Optional) Enter the backup RADIUS server to use if the primary server does not respond. If the server returns an authentication failure (bad password, unknown username, etc), then the backup server is not tried. The backup server is tried only when the primary server has a communication failure. Use the same format as the primary server to specify the host name and port.
Shared Authentication Secret	Enter a RADIUS passphrase.
NAS IP Address	The IP address of the Network Access Server (NAS).
Request Timeout	The length of time, in seconds, to wait for a response from the RADIUS server (in seconds). The default is 10 .
Retry Request	Number of times to retry a RADIUS request. The default is 1 .
RADIUS Protocol:	Use the drop-down menu to choose a protocol option. The default is None .

- 7 Click **Save**.

Local Password Fallback

You can use this feature to log in using your local user name and password if the external authentication (Certificate, LDAP, or RADIUS) fails, if you forgot your password to the authentication server, or if the authentication server is not available.

The Use Local Authentication allows the default admin to log in even when the remote authentication server is not available, by adding a **Use Local Authentication** checkbox to the login screen. Out-of-box, this option is enabled only for the default administrator. However, it is possible to allow local password fallback for all users. For example, you could configure the RADIUS authentication method to allow users to log in using local authentication instead of RADIUS should they fail to authenticate to the configured external RADIUS server(s).

For information on how to allow local password fallback for all users for all users, see [“Client Certificate Authentication” on page 397](#), [“LDAP/AD and LDAPS Authentication” on page 398](#), or [“RADIUS Authentication” on page 400](#).

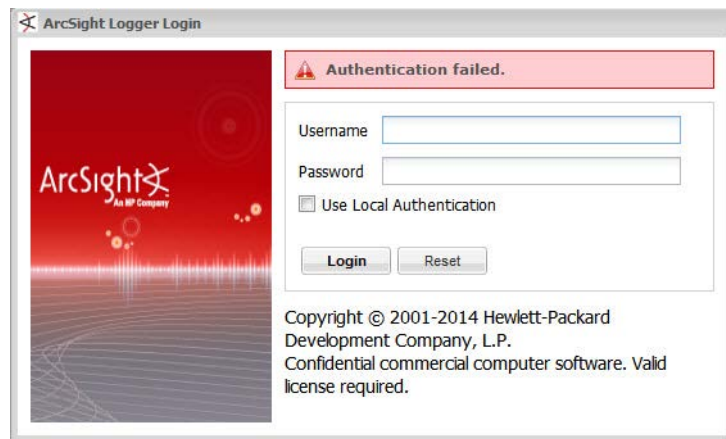
To log in when authentication fails:

- 1 Mark the **Use Local Authentication** checkbox if the login failure was caused by failure of the external authentication.



Note

This option is only available to the default admin unless it has been enabled for other users.

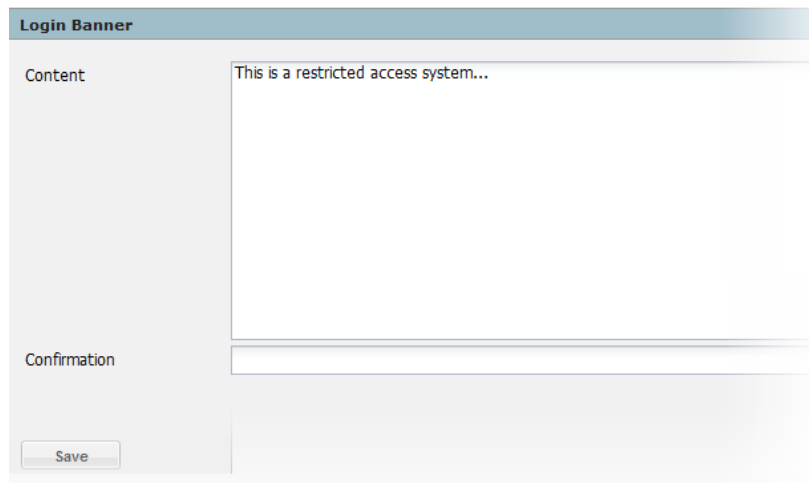


- 2 Enter your login and password and click **Login**.

Login Banner

You can customize the message on the login screen to suit your needs. The text you enter in the Content field is displayed above the Username and Password fields on the login screen. In addition, you can enter a confirmation message that the user must click to enable the Username and Password fields.

Login Banner



You must have the “Configure Login Settings” permission enabled for your user account to edit the login banner.

To customize the login banner:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Login Banner** in the **Users/Groups** section.
- 3 Enter the text you want to display as the login banner in the Content field.

You can enter only unformatted text in this field; however, you can apply standard HTML tags to display formatted text. Loading images in this field is not allowed.

- 4 (Optional) Enter text in the Confirmation field.

If you enter text in this field, the text will be accompanied by a checkbox that the user must click to enable the Username and Password fields. For example, if you enter “Are you sure?”, “Do you want to proceed?”, or “I agree.” in this field, the user must click the checkbox in order to log in.

- 5 Click **Save**.

User Management

The **Users** and **Groups** tabs enable you to manage users and user groups on your system. User groups are a way to enforce access control to various sections of your system.

Users

Open the **Users** tab to manage the users that can log in to your system. You can add a new user, edit user information, or delete a user at any time. You must have the appropriate System Admin group rights to perform these functions.

To add a new user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 In the **Users** tab, click **Add** from the top left side of the page.

4 Enter the following parameters.

Parameter	Description
<i>Credentials</i>	
Login	The user's login name.
Password	The user's password.
Confirm Password	Reenter the users' password.
<i>Contact Information</i>	
Use Client DN	<p>If you enabled SSL client certificate or LDAP authentication, click this link to enter user's the Distinguished Name (Certificate Subject) information. The Distinguished Name should be similar to this format:</p> <p>CN=UserA,OU=Engg Team,O=ArcSight\, Inc.,L=Cupertino,C=US,ST=California</p> <p>To determine the DN, use this URL to display the certificate:</p> <p><code>https://<hostname or IP address>/platform-service/DisplayCertificate</code></p> <p>OR</p> <p>Obtain the DN information for a user from the browser that the user will open to connect to the system. For example, on Firefox, click Tools > Options > Advanced > Encryption > View Certificates > Your Certificates > Select the certificate > View.</p>
First Name	The user's first name.
Last Name	The user's last name.
Email	The user's email address.
Phone Number	(Optional) The user's phone number.
Title	(Optional) The user's title.
Department	(Optional) The user's department.
Fax	(Optional) The user's fax number.
Alternate Number	(Optional) The user's alternate phone number.
Assign to Groups	Select the groups to which this user belongs. This setting controls the privileges a user has on this Logger.
Notes	(Optional) Other information about the user.

5 Click **Save and Close**.

To edit a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 In the **Users** tab, select the user (or users) you want to edit.
- 4 Click **Edit** from the top left side of the page.

5 Update the user information as necessary.

6 Click **Save User**.

To delete a user:

1 Click **System Admin** from the top-level menu bar.

2 Click **User Management** in the **Users/Groups** section in the left panel.

3 In the **Users** tab, select the user (or users) you want to delete.

4 Click **Delete** from the top left side of the page.

Reset Password

The Reset Password feature enables you to reset a user's password without knowing their password. If you are using an SMTP-configured server and have permissions to create and update users, you can reset a user's password by clicking the **Reset Password** button. An automated email is sent to the user with the new password string.

An SMTP server must be configured for the automated email containing the temporary password to be sent. If an SMTP server is not configured, the password will not be reset because an email cannot be sent.

To reset a user's password:

1 Click **System Admin** from the top-level menu bar.

2 Click **User Management** in the **Users/Groups** section in the left panel.

3 In the **Users** tab, select the user (or users) whose passwords you want to reset.

4 Click **Reset Password** from the top left side of the page.

The user must use the temporary string to log in within the time specified in the email. If the user does not log in within the specified time, the account becomes deactivated. If the account has been deactivated, the admin must re-activate it before resetting the password.

To activate a user:

1 Click **System Admin** from the top-level menu bar.

2 Click **User Management** in the **Users/Groups** section in the left panel.

3 In the **Users** tab, select the user (or users) that you want to activate.

4 Choose **Edit**.

5 Check the **Active** box.

6 **Save** the changes.

Groups

User groups define privileges to specific functions on your system and serve to enforce access control to these functions. For example, if you want User A to be able to run searches but not reports, assign that user to the Search group but not to the Reports group.

User groups are organized by the following types: System Admin, Logger Rights, Logger Search, Logger Reports, and Connector Appliance Rights (only on Logger platforms with integrated Connector Appliance). Each type has a pre-defined, default user group in which all privileges for the type are enabled. To authorize a subset of the privileges for a specific

group type, create a new user group and enable only the privileges you want to provide for that group. Then, assign restricted users to the newly created group.

System Admin Groups

The System Admin Group controls the system administration operations for your system, such as configuring network information, setting storage mounts, installing SSL certificates, and user management.

Read Only System Admin Group

In addition to the default System Admin Group that enables all rights (privileges), a Read Only System Admin Group is available on your system. Users assigned to this group can view System Admin settings, but cannot change them.

Refer to your system's user interface for a complete list of privileges available to this group.

Logger Rights Group

The Logger Rights Group controls the Logger application operations for your system, such as viewing the Logger dashboards and configuring all the settings in the Configuration menu (including event archives, storage groups, alerts, filters, and scheduling tasks.)

Refer to your system's user interface for a complete list of privileges available to this group.

Logger Search

The Logger Search Group controls local and peer searches through the following privileges:

- Search for events
- Search for events on remote peers

If the group is configured to allow users to run local and peer searches, users assigned to this group can perform those operations. Conversely, if the group is configured to prevent users from running local and peer searches, users assigned to this group cannot perform those operations.

Logger Reports

The Logger Reports group controls all report operations on Logger such as run, edit, delete, schedule, and view published reports.

Refer to your system's user interface for a complete list of privileges available to this group.

Connector Appliance Rights Groups

The Connector Appliance Rights Group controls the Connector Appliance application operations for your system, such as viewing the Connector Appliance dashboards and backup operations.

Read Only Connector Appliance Group

In addition to the Default Connector Appliance Rights Group that enables all privileges to the Connector Appliance functions, a Read Only Connector Appliance Group is available on your system. A read-only user can view the tabs and the operations displayed on the tabs, and can perform operations such as refresh, view certificate list, and logfu.

Refer to your system's user interface for a complete list of rights available to this group.




In the **Default Connector Appliance Group**, under Application Options:

- Option rights can now be found here, not under System Admin Groups.

Two new management rights (or privileges) allow for control of the **Manage** and **Repositories** tabs.

Managing a User Group

To create a new user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 Click the **Groups** tab.
- 4 Click **Add** from the top left side of the page.
- 5 Define the new group:
 - a In the **Group Name** field, provide a name for the group.
 - b In the **Description** field, provide a description for the group.
 - c From the Group Type drop-down box, select the group type.
 - d Click the down arrow icon () next to the group type name to view and select privileges that you want to assign to the users in this group.
- 6 Click **Save and Close** to save the settings of the group, or click **Save and Edit Membership** to add users to this group.

To edit a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 Click the **Groups** tab.
- 4 Select the group that you want to edit, and click **Edit** at the top left side of the page.
- 5 Update the user group information.

If you need to edit the group's membership:

- a Click **Save and Edit Membership** to display the Edit Group Membership page.
- b Click **Add** from the top left of the Edit Group Membership page.
- c Select users you want to add. By default, you can add only users who do not belong to other groups of the type that you are editing. To add such users, click **Show users that belong to other <group_type> groups**.

When you add a user who belongs to another group of the same group type as the one you are updating, that user is automatically removed from the previous group.

- d Click **OK**.
 - e Click **Back to Group List**.
- 6 Click **Save and Close**.

To delete a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 Click the **Groups** tab.
- 4 Select the group (or groups) that you want to delete.
- 5 Click **Delete** at the top left side of the page.

Change Password

You can use the **Change Password** menu to change your password. This feature is available to all users for changing their passwords, unlike the Reset Password feature that enables a system administrator to reset the password of users without knowing the password. Passwords are subject to the password policy specified by the Admin user.

To change your password:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Change Password** in the **Users/Groups** section in the left panel to display the Change Password for <User Name> page.
- 3 Enter the Old Password, the New Password, and enter the New Password a second time to confirm.
- 4 Click **Change Password**.

Other System Administration Information

This section contains information related to system administration that you will need to fully administer your Logger, including system health events and SNMP polling.

Monitoring System Health

You can monitor your Logger's health in these ways:

- By using a pre-defined system filter, as listed in ["System Filters/Predefined Filters" on page 128](#). The pre-defined system health filters are based on the system health events listed in ["System Health Events" on page 408](#).
- By searching for system health events in Logger's Internal Storage Group, as listed in ["System Health Events" on page 408](#). If a pre-defined system health filter does not suit your needs, you can create alerts based on the system health events.
- By polling system health events, as explained in ["Polling System Health Information Using SNMP" on page 366](#). You can poll system health information from your system by using SNMP version 2 from any standard network management system.

To set up notification of system health events:

- 1 Configure the Logger's SMTP settings (see ["SMTP" on page 363](#)) or create an SNMP Destination (see ["Sending Notifications to SNMP Destinations" on page 307](#)) or Syslog Destination (see ["Sending Notifications to Syslog Destinations" on page 308](#)).
- 2 Create an Alert that uses one or more System Alert Filters or define a query that searches for the system health events in Logger's Internal Storage Group, and specify match count and threshold (see ["Alerts" on page 296](#)).
- 3 Enable the new Alert.

System Health Events

The following table lists the system health events that Logger generates. These events are also referred as Logger Internal Events because they are stored in Logger's Internal Storage Group. See [Appendix F, Examples of System Health Events on page 151](#) for examples of these events.

The pre-defined System Filters that provide system health status are based on some of these events. If a pre-defined filter does not suit your needs, create an alert using one of these events.

Starting with Logger 5.1, the format in which system health events are generated was changed to provide more meaningful information. These changes include:

- Addition of new events (for example, Current and Voltage).
- Instead of referring to all system health events as Logger Internal Event in the `name` field, meaningful names are used (for example, Fan OK, Temperature OK).
- Three severity levels for each event have been added to the `agentSeverity` field—1 (OK), 5 (Degraded), and 8 (Severe).
- The `deviceCustomString` and `deviceCustomStringLabel` field mappings have changed. Refer to a specific event to see the changes.
- Device Event Class ID (`deviceEventClassId`) and Device Event Category (`deviceEventCategory`) of the events have changed. An updated list is available in the following table.
- All hardware-related events are classified as `hardware:nnn` events, where `nnn` is a three-digit number that identifies the hardware component (for example, `hardware:13x` identifies the fan events.)

If you are upgrading from Logger 5.0 Patch 2 or earlier, any existing filters or queries based on earlier events will not work on the events collected after the upgrade. Those filters and queries will continue to work on the events collected before the upgrade. The pre-defined System Filters are compatible with both the new and the old formats.



- The sensor names in each event are hardware specific; therefore, they are not consistent across various Logger platforms. Use the event name (Name) and status (CustomString3) fields to determine the status of a sensor. The raw status (CustomString4), location (CustomString5), and sensor name (CustomString6) fields are for informational use when diagnosing a hardware problem and are not consistent across appliance types.
- HP recommends that you develop custom alerts for certain System Health Events to prevent users from being alerted too often. Some of the conditions that your system alerts on may be self-clearing or warnings that you do not want to be alerted about until a specific number of warnings have been generated.

Group	Device Event Category	Device Event Class ID
System Health Events for appliance and software Loggers		
CPU	/Monitor/CPU/Usage	cpu: 100
Disk	/Monitor/Disk/Read	disk: 102
	/Monitor/Disk/Write	disk: 103

Group	Device Event Category	Device Event Class ID
EPS	/Monitor/Receiver/EPS/All	eps: 100
	/Monitor/Receiver/EPS/Individual	eps: 102
	/Monitor/Forwarder/EPS/All	eps: 101
	/Monitor/Forwarder/EPS/Individual	eps: 103
Memory	/Monitor/Memory/Usage/Platform	memory: 100
Network	/Monitor/Network/Usage/In	network: 100
	/Monitor/Network/Usage/Out	network: 101
Search	/Monitor/Search/Performed	search: 100
Storage Group	/Monitor/StorageGroup/Space/Used	storagegroup: 100
	Note: The size of the storage group, indicated by the “fsize” field is in GB.	
System Health Events for appliance Loggers only		
Battery	/Monitor/Sensor/Battery/OK	hardware: 121
	/Monitor/Sensor/Battery/Degraded	hardware: 122
	/Monitor/Sensor/Battery/Failed	hardware: 123
Current (Electrical)	/Monitor/Sensor/Current/OK	hardware: 101
	/Monitor/Sensor/Current/Degraded	hardware: 102
	/Monitor/Sensor/Current/Failed	hardware: 103
Disk	/Monitor/Disk/Space/Remaining/Root	disk: 101
Fan	/Monitor/Sensor/Fan/OK	hardware: 131
	/Monitor/Sensor/Fan/Degraded	hardware: 132
	/Monitor/Sensor/Fan/Failed	hardware: 133
Power Supply	/Monitor/Sensor/PowerSupply/OK	hardware: 141
	/Monitor/Sensor/PowerSupply/Degraded	hardware: 142
	/Monitor/Sensor/PowerSupply/Failed	hardware: 143
RAID	/Monitor/RAID/Controller/OK	raid: 101
	/Monitor/RAID/Controller/Degraded	raid: 102
	/Monitor/RAID/Controller/Failed	raid: 103
	/Monitor/RAID/BBU/OK	raid: 111
	/Monitor/RAID/BBU/Degraded	raid: 112
	/Monitor/RAID/BBU/Failed	raid: 113
	/Monitor/RAID/Disk/OK	raid: 121
	/Monitor/RAID/Disk/Rebuilding	raid: 122
	/Monitor/RAID/Disk/Failed	raid: 123

Group	Device Event Category	Device Event Class ID
Temperature	/Monitor/Temperature/OK	hardware: 151
	/Monitor/Temperature/Degraded	hardware: 152
	/Monitor/Temperature/Failed	hardware: 153
Voltage	/Monitor/Sensor/Voltage/OK	hardware: 111
	/Monitor/Sensor/Voltage/Degraded	hardware: 112
	/Monitor/Sensor/Voltage/Failed	hardware: 113

Using the Command Line Interface

To use the Command Line Interface (CLI), attach a terminal to the serial port on Logger or attach a monitor and keyboard. The serial port expects a standard VT100-compatible terminal: 9600 bps, 8-bits, no parity, 1 stop bit (8N1), no flow control.

The following commands are available at the CLI prompt:

Category	Command	Description
System Commands	exit	Logout
	halt	Stop and power down the Logger appliance
	help	Opens the command line interface help
	reboot	Reboot the Logger appliance
Admin	show admin	Show the default administrator user's name
Authentication	reset authentication	Reset to local authentication
Config	show config	Show host name, IP address, DNS, and default gateway for the Logger
Date	show date	Show the date and time currently configured on the Logger
	set date	Set the date and time on Logger. The date/time format is yyyyMMddmmhhss. Example date: 20101219081533
Default Gateway	set defaultgw <IP> [nic]	Set the default gateway for one or all network interfaces
	show defaultgw [nic]	Display the default gateway for all or the specified network interface
DNS	show dns	Show the currently configured DNS servers on the Logger

Category	Command	Description
	set dns <sd> <ns> set dns <sd1>,<sd2> <ns1> <ns2>	Set DNS name server(s). sd=search domain, ns = name server You can add up to three name servers and six search domains. Note: When using multiple search domains, separate them with a comma, but no space. When using multiple name servers separate them with a space but no comma.
Hostname	show hostname set hostname <host>	Show the currently configured hostname on the Logger Set Logger's host name
IP	show ip [nic] set ip <nic> <IP> [/prefix] [netmask]	Show the IP addresses of all or the specified network interface Set Logger's IP address for a specific network interface
NTP	set ntp <ntp server> <ntp server> <ntp server> ... show ntp	Sets the NTP server addresses. This entry overwrites the current NTP server setting. You can specify as many NTP servers as you like. If you specify multiple NTP servers, they are each checked in turn. The time given by the first server to respond is used. Example: logger> set ntp ntp.arcsight.com time.nist.gov 0.rhel.pool.org Show the current NTP server setting. Example: logger> show ntp ntp.arcsight.com time.nist.gov 0.rhel.pool.org
Password	set password	Set the password the current user's account
Process	restart process start process status process stop process	Restart a process Start a process Show process status Stop a process
SSL Certificate	show sslcert reset sslcert restart sslcert diag sslcert	Show the currently loaded SSL certificate on Logger Install and restart the HTTPS server with the default, self-signed certificate that Logger shipped with Restart the HTTPS server Display the SSL session information
Status	show status	Show the Logger configuration

Chapter 8

System Admin - Software Logger

This chapter describes the System Administration tools that enable you to create and manage users and user groups, and to configure SMTP and other system settings.

This chapter includes information on the following areas of system administration for the Software Logger

- ["System Locale" on page 413](#)
- ["SMTP" on page 414](#)
- ["License & Update" on page 414](#)
- ["Process Status" on page 415](#)
- ["System Settings" on page 416](#)
- ["Audit Logs" on page 416](#)
- ["Audit Forwarding" on page 417](#)
- ["SSL Server Certificate" on page 417](#)
- ["SSL Client Authentication" on page 422](#)
- ["FIPS 140-2" on page 424](#)
- ["Authentication" on page 428](#)
- ["Login Banner" on page 437](#)
- ["User Management" on page 438](#)
- ["Change Password" on page 442](#)
- ["Monitoring System Health" on page 443](#)

System

From the System tab, you can configure system specific settings such as network settings (if applicable) and SMTP.

System Locale

The System Locale setting ensures that the user interface displays information such as date, time, numbers, and messages in the format and language appropriate for the selected country.

The System Locale is configured during the Logger installation process. Once configured it cannot be changed.

To view the System Locale:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **System Locale** in the **System** section.
The System Locale Setting dialog box displays the Locale.

SMTP

Your system uses the Simple Mail Transfer Protocol (SMTP) setting to send email notifications such as alerts and password reset emails.

To add or change SMTP settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SMTP** in the **System** section and enter these settings.

Setting	Description
Primary SMTP Server	The IP address or hostname of the SMTP server that will process outgoing email.
Backup SMTP Server	The IP address or hostname of the SMTP server that will process outgoing email in case the primary SMTP server is unavailable.
Outgoing Email Address	The email address that will appear in the From: field of outbound email.

- 3 Click **Save**.

**Note**

Be sure to configure your reports to use the same SMTP settings. For instructions, see [“Report Server Administration” on page 238](#).

License & Update

This page displays license information, the version of the components, and the elapsed time since Logger was last restarted. To view details of your license, open **Configuration** from the top-level menu bar, and then click **License Information**. For details, see [“License Information” on page 351](#).

Updating the License File

To update a license file:

- 1 Download the update file from the HP Customer Support site (SSO) at <https://support.openview.hp.com> to the computer from which you can connect to the Logger with your browser.

For more information, see [“Acquiring a License for a Software Logger” on page 41](#).

- 2 From the computer to which you downloaded the update file, log in to the Logger user interface using an account with administrator (upgrade) privileges.
- 3 Click **System Admin** from the top-level menu bar.
- 4 Click **License & Update** in the **System** section.
- 5 Browse to the license file you downloaded earlier, and click **Upload Update**.

An “Update In Progress” page displays the update progress.

Once the update has completed, the Update Results page displays the update result (success/failure). If you are only installing or updating a license, a restart is not required.


Process Status

The **Process Status** page lists all processes related to your system and enables you to view the details of those processes and start, stop, or restart them.

To view the Process Status page:







- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Process Status** in the **System** section to display a page similar to the one shown in the following figure.

Process Status


 Refresh Status

System						
System	Status	Load	CPU Usage	Memory Usage	Data Collected	
n035-h027.qa.arcsight.co	running	[1.89] [1.65] [1.89]	16.8%us 1.9%sy 1.2%wa	68.0% [4076648 kB]	09/15/2010 15:01:23	


NOTE: This Start/Stop buttons are for diagnostic purposes. Please use them with care.

Processes				
 Start  Stop  Restart				
Process	Status	Uptime	CPU Usage	Memory Usage
 apache	running	4h 25m	0.0%	0.0% [3320 kB]
 aps	running	4h 25m	0.3%	5.0% [303676 kB]
 connector	running	4h 15m	0.0%	0.0% [608 kB]

In the **Processes > Process** list on this screen, “processors” refers to forwarders.







- 3 To view the details of a process, click the  icon to the left of the process name, as shown in the following figure.

Process Status

 Refresh Status

System						
System	Status	Load	CPU Usage	Memory Usage	Data Collected	
n035-h027.qa.arcsight.co	running	[1.50] [1.84] [1.90]	3.9%us 0.5%sy 2.1%wa	68.4% [4097992 kB]	09/15/2010 15:07:54	

NOTE: This Start/Stop buttons are for diagnostic purposes. Please use them with care.

Processes				
 Start  Stop  Restart				
Process	Status	Uptime	CPU Usage	Memory Usage
 apache	running	4h 32m	0.0%	0.0% [3320 kB]
Children	0			
CPU Percent	0.0%			
CPU Percent Total	0.0%			
Data Collected	09/15/2010 15:08:09			
Memory Kilobytes	3320			
Memory Kilobytes Total	96128			
Memory Percent	0.0%			
Memory Percent Total	1.6%			
Monitoring Status	monitored			
Parent PID	1			
PID	28151			
Status	running			
Uptime	4h 32m			
 aps	running	4h 32m	0.2%	5.1% [311040 kB]
 connector	running	4h 22m	0.0%	0.0% [608 kB]

To start, stop, or restart a process, select the process and click **Start**, **Stop**, or **Restart** at the top of the process list.

System Settings

If you did not select Logger to start as service during the installation process, you can do so using the **System Settings** page. When you select this option Logger will use a service called `arcsight_logger`, enabled to run at levels 2, 3, 4, and 5.

To configure Logger to start as a service:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **System Settings** in the left panel.
- 3 From under **Service Settings**, choose the appropriate option:
 - ◆ Start as a Service
 - ◆ Do not start as a Service
- 4 Click **Save**.

Logs

Your system can generate audit logs at the application and platform levels. Use the Logs sub-menu to search audit logs.

Audit Logs

Your system's audit logs are available for viewing. Audit logs, as Common Event Format (CEF) audit events, can be sent to ArcSight ESM directly for analysis and correlation. For information about forwarding audit events, see [“Audit Forwarding” on page 417](#).

Audit Logs

Search Audit Logs

Timestamp

12/04/2012

12/04/2012

Description

User

Search

Search Results

User	Description	Timestamp
admin	Successful login	2012/12/04 13:00:21
admin	Authentication Session settings successfully changed.	2012/12/04 11:44:06
admin	Successful login	2012/12/04 11:43:51
admin	Saved Search [SL_Saved_Search_1204_1] has been added	2012/12/04 11:28:06
admin	Saved Search [SL_Saved_Search_1204_2] has been added	2012/12/04 11:28:06
admin	Dashboard [SL_Dashboard_1204_1] has been added	2012/12/04 11:28:06

To view audit logs:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit Logs** in the **Logs** section.
- 3 Select the date and time range for which you want to obtain the log.

- 4 (Optional) To refine the audit log search, specify a string in the Description field and a user name in the User field. When a string is specified, only logs whose Description field contains the string are displayed. Similarly, when a user is specified, only logs whose User field contains the username are displayed.
- 5 Click **Search**.

Audit Forwarding

You can forward audit events to an ArcSight ESM for correlation and analysis. For a list of audit events that you can forward, see [Appendix C, Logger Audit Events, on page 559](#).

To forward audit events to specific ESM destinations:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Audit Forwarding** in the **Logs** section.
- 3 Select destinations from the **Available Destinations** list and click the right arrow icon (➡) to move the selected destination to the **Selected Destinations** list.

You can select multiple destinations at the same time and move them, or you can move all available destinations by clicking the (➡) icon.

Audit Forwarding

The screenshot shows a window titled "Audit Destinations". It contains two main sections: "Available Destinations" on the left and "Selected Destinations" on the right. In the "Available Destinations" list, the item "test1" is selected. Between the two lists are four arrow icons: a right-pointing arrow, a double right-pointing arrow, a left-pointing arrow, and a double left-pointing arrow. At the bottom left of the window is a "Save" button.

The destinations are ESM destinations that you configure on the ESM Destinations page (**Configuration** (or **Configuration > Settings**) > **Event Input/Event Output > ESM Destinations**).

- 4 Click **Save Settings**.

Security

Security settings enable you to configure SSL server certificates, enable and disable FIPS (Federal Information Processing Standards) mode on your system, and configure SSL client authentication for client certificate and Common Access Card (CAC) support.



For steps on how to create a user DN, see ["Users" on page 438](#), and refer to the section "Use Client DN" in the parameters table.

SSL Server Certificate

Your system uses Secure Sockets Layer (SSL) technology to communicate securely over an encrypted channel with its clients, such as SmartConnectors, when using the

SmartMessaging technology and other ArcSight systems. Your system ships with a self-signed certificate so that an SSL session can be established the first time you use the appliance. For more information on this option, see [“Generating a Self-Signed Certificate” on page 418](#).

Although a self-signed certificate is provided for your use, HP strongly recommends using a certificate authority (CA) signed certificate. Even if FIPS is not enabled on your system, it must use a **CA-signed certificate** if it is a destination of a FIPS-enabled container (on a Connector Appliance) or a software-based SmartConnector. Additionally, ensure that the root certificate of the CA that signed your system's certificate is trusted on the container or the SmartConnector. If the CA's root certificate is not trusted, load it on the container by following instructions in the “Managing Certificates on a Container” section in the Connector Appliance Administrator's Guide.

To facilitate obtaining a CA-signed certificate, your system can generate a Certificate Signing Request. Once a signed certificate file is available from the CA, it can be uploaded to your system for use in a subsequent authentication. For detailed instructions, see [“Generating a Certificate Signing Request \(CSR\)” on page 419](#).

Your system generates an audit event when the installed SSL certificate is going to expire in less than 30 days or has already expired. The event with Device Event Class ID “platform:407” is generated periodically until you replace the certificate with one that is not due to expire within 30 days.

Generating a Self-Signed Certificate

Your appliance ships with a self-signed certificate so that an SSL session can be established the first time you connect. This type of certificate does not require signing from another entity and can be used immediately.

To generate a self-signed certificate:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** from the **Security** section in the left panel to display the **Generate Certificate/Certificate Signing Request** page.
- 3 Click the **Generate Certificate** tab.

SSL Settings

Generate Certificate Import Certificate

Generate Certificate/ Certificate Signing Request

Enter Certificate Settings

Country (2-letter code)	US
State/Province	California
City/Locality	Sunnyvale
Organization Name	Hewlett-Packard
Organizational Unit	Support Team
Hostname	scsm0102.hpwelbts.adagga.hp.com
Email Address	arst-support@hp.com
Private Key Length	1024

Generate CSR Generate Certificate View Certificate

- 4 From the **Enter Certificate Settings** field, enter new values for the following fields:

Parameter	Description
Country	A two-letter country code, such as 'US' for the United States.
State/Province	State or province name, such as 'California.'
City/Locality	City name, such as 'Sunnyvale'.
Organization Name	Company name, governmental entity, or similar overall organization.
Organizational Unit	Division or department within the organization.
Hostname	The host name or IP address of this system. When specifying the host name, make sure that this name matches the name registered in the Domain Name Service (DNS) server for the system. Note: If the host name or IP address of this system changes in the future, you must generate a new self-signed certificate or CSR. Once a new certificate is obtained, you must upload it to ensure that the connectors (in FIPS mode) which communicate with the system are able to validate the host name.
Email Address	The email address of the administrator or contact person for this CSR.
Private Key Length	Select the length (in bits) of the private key: 1024 , 2048 , 4096 , or 8192 .

Use the first two buttons to generate a CSR or a self-signed certificate. The **View Certificate** button is only used to view the resulting certificate.

Button	Description
Generate CSR	Click to generate a Certificate Signing Request (CSR).
Generate Certificate	Click to generate a self-signed certificate.
View Certificate	Click to view the generated certificate.

- 5 Click the **Generate Certificate** button to generate the self-signed certificate.



The Apache server restarts while generating the certificate. You may get an error communicating to the web server while this is happening. This is expected behavior, and communication is automatically restored once Apache is back up.

- 6 Click **Ok** after the confirmation message appears.
- 7 Click the **View Certificate** button to view the PEM encoded self-signed certificate.

Generating a Certificate Signing Request (CSR)

The first step in obtaining a CA-signed certificate is to generate a Certificate Signing Request (CSR). The CSR must be generated on the system for which you are requesting a

certificate. That is, you cannot generate a CSR for System A on System B or use a third-party utility for generation.

The resulting CSR must be sent to a CA, such as VeriSign, which responds with a signed certificate file.

To generate a certificate signing request:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** from the **Security** section in the left panel to display the **Generate Certificate/Certificate Signing Request** page.
- 3 Click the **Generate Certificate** tab.

SSL Settings

Generate Certificate Import Certificate

Generate Certificate/ Certificate Signing Request

Enter Certificate Settings

Country (2-letter code)	US
State/Province	California
City/Locality	Sunnyvale
Organization Name	Hewlett-Packard
Organizational Unit	Support Team
Hostname	scm01c.hpsw-lbns.adappa.hp.com
Email Address	arst-support@hp.com
Private Key Length	1024

Generate CSR Generate Certificate View Certificate

- 4 From the **Enter Certificate Settings** field, enter new values for the following fields:

Parameter	Description
Country	A two-letter country code, such as 'US' for the United States.
State / Province	State or province name, such as 'California.'
City / Locality	City name, such as 'Sunnyvale'.
Organization Name	Company name, governmental entity, or similar overall organization.
Organizational Unit	Division or department within the organization.
Hostname	<p>The host name or IP address of this system.</p> <p>When specifying the host name, make sure that this name matches the name registered in the Domain Name Service (DNS) server for the system.</p> <p>Note: If the host name or IP address of this system changes in the future, you must generate a new self-signed certificate or CSR. Once a new certificate is obtained, you must upload it to ensure that the connectors (in FIPS mode) which communicate with the system are able to validate the host name.</p>

Parameter	Description
Email Address	The email address of the administrator or contact person for this CSR.
Private Key Length	Select the length (in bits) of the private key: 1024 , 2048 , 4096 , or 8192 .

Use the first two buttons to generate a CSR or a self-signed certificate. The **View Certificate** button is only used to view the resulting certificate.

Button	Description
Generate CSR	Click to generate a Certificate Signing Request (CSR).
Generate Certificate	Click to generate a self-signed certificate.
View Certificate	Click to view the generated certificate.

- 5 Choose **Generate CSR** to generate a certificate signing request.
- 6 If the CSR was successfully generated, a pop-up window appears, allowing you to either download the CSR file or to cut-and-paste its content.

To do so, copy all the lines from -----BEGIN CERTIFICATE REQUEST----- to -----END CERTIFICATE REQUEST-----.
- 7 Send the CSR file to your certificate authority to obtain the CA-signed certificate.
- 8 Once the CA-signed certificate file is obtained, continue on to [Importing a Certificate](#) below.

Importing a Certificate

If you have obtained a certificate from your certificate authority (CA), follow the steps below to import it onto your system.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Server Certificate** under the **Security** section in the left panel.
- 3 Select the **Import Certificate** tab.

SSL Settings

Generate Certificate
Import Certificate

Import Signed Certificate

Note: After uploading the new certificate, close and re-open the browser.

Select the Signed Certificate File

Last Import Status

Last SSL Certificate Import Status

--- No Status Exists ---

- 4 Click the **Browse** button to locate the signed certificate file on your local file system.



The imported certificate must be in **Privacy Enhanced Mail (PEM)** format.

- 5 Click **Import and Install** to import the specified certificate.
- 6 If using **HTTPS** and depending on your browser, you may need to close and restart the browser for the new certificate to take effect. If you are unsure of your browser's requirements, close and restart it.

SSL Client Authentication

Your system supports client authentication using SSL certificates. SSL client authentication is a form of two-factor authentication that can be used *as an alternate* or *in addition to* local password authentication. As a result, your system can be configured for SmartCards, such as Common Access Card (CAC) based authentication. CAC is a standard identification card for active duty members of the Uniformed Services, Selected Reserve, DOD civilian employees, and eligible contractor personnel.



CAC is a form of client certificate authentication. Information on client certificate authentication applies to CAC.

Your system also supports LDAPS authentication. The SSL certificate for the LDAPS server must be uploaded into the trusted store. After uploading the SSL certificate, the aps process must be restarted (**System Admin > Process Status > aps > Restart**).

Configuring Logger to Support SSL Client Authentication

Perform the following steps to configure Logger to support SSL client authentication.

On the Logger:

- 1 If the Logger uses the default signed certificate it shipped with from ArcSight, replace it with a *FIPS-compliant*, signed SSL server certificate. Follow instructions at ["Uploading Trusted Certificates" on page 423](#) to load the certificate.



All SSL client certificates used for authentication must be FIPS compliant (that is, hashed with FIPS-compliant algorithms) even if FIPS is not enabled on your Logger.

- 2 Enable client certificate authentication, as described in ["Client Certificate Authentication" on page 433](#).
- 3 If the client certificates are **CA signed**, upload the root certificate of the authority who signed the certificates that will be used for authenticating clients, as described in ["Uploading Trusted Certificates" on page 423](#).

If the client certificates used to authenticate with Logger are signed by different CAs, make sure you upload root certificates of **all** CAs.

If the client certificates are **self-signed**, upload the public portion of the client certificate.

- 4 Configure a user name for each user who will be connecting to the Logger using a client certificate, as described in [“User Management” on page 438](#).
- 5 (Optional) Upload a certificate revocation list (CRL), as described in [“Uploading a Certificate Revocation List” on page 424](#).
- 6 (Optional) If this Logger is configured to use **only** SSL Client Authentication, make sure this Logger's Authorization ID and Code are appropriately configured on other Loggers that with it. For more information, see [“Peer Loggers” on page 326](#).

On the Client (Web browser):

Configure your browser to provide the SSL client certificate when accessing Logger.
(Upload the private key in PKCS 12 format in your browser.)

Uploading Trusted Certificates

A trusted certificate is used to authenticate users that log in to your system. Uploading a trusted certificate is required if you are using LDAPS authentication. The trusted certificate is used to authenticate the remote LDAPS server. The certificate needs to be in Privacy Enhanced Mail (PEM) format.

To upload a trusted certificate:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Client Authentication** in the **Security** section in the left panel.
- 3 On the **Trusted Certificates** tab, click **Browse** to find the trusted certificate on your local file system.
- 4 Click **Upload**.

The trusted certificate is uploaded and listed in the “Certificates in Repository” list on the same page where you uploaded it.

The screenshot shows the 'Trusted Certificates' tab in the software interface. It includes an 'Upload Certificate' section with a message indicating the certificate must be in PEM format, an 'Upload File:' label, a text input field, a 'Browse...' button, and an 'Upload' button. Below this is a 'Certificates in Repository' section with a table. The table has columns for 'Certificate Name', 'Start date', and 'Expiration Date'. The 'Certificate Name' column contains a checkbox and a link. A 'Delete' button is located at the bottom of the table.

To view details about a trusted certificate, click the link displayed in the Certificate Name column.

To delete a trusted certificate, select the certificate and click **Delete**.

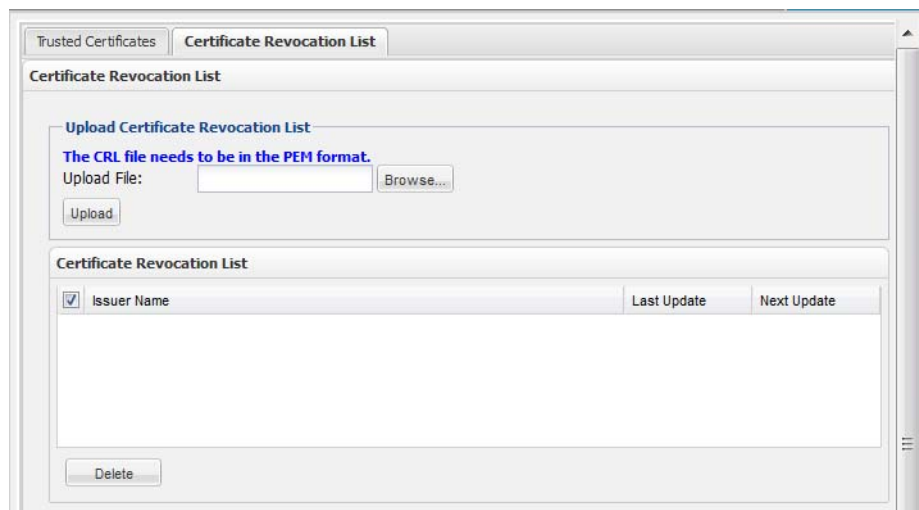
Uploading a Certificate Revocation List

A certificate revocation list (CRL) is a computer-generated record that identifies certificates that have been revoked or suspended before their expiration dates. To support CAC, you need to upload a CRL file to your ArcSight system. The CRL file needs to be in PEM format.

To upload a CRL file:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **SSL Client Authentication** in the **Security** section in the left panel.
- 3 In the **Certificate Revocation List** tab, click **Browse** to find the CRL file on your local file system.
- 4 Click **Upload**.

The CRL is uploaded and listed in the Certificate Revocation List.



To view details about a CRL, click the link displayed in the Issuer Name column.

To delete a CRL file, select it and click the **Delete** button.

Enabling Client Certificate Authentication

To enable client certificate authentication, see [“Client Certificate Authentication” on page 433](#).

FIPS 140-2

Your system supports the Federal Information Processing Standard 140-2 (FIPS 140-2). FIPS 140-2 is a standard published by the National Institute of Standards and Technology (NIST) and is used to accredit cryptographic modules in software components. The US Federal government requires that all IT products dealing with Sensitive, but Unclassified (SBU) information meet these standards.

If your system needs to be FIPS 140-2 compliant, you can enable FIPS. Once you do so, the system uses the cryptographic algorithms defined by the NIST for FIPS 140-2 for all encrypted communication between its internal and external components.



To be fully FIPS 140-2 compliant, all components of your Logger deployment need to be in FIPS 140-2 mode. For example, if you enable FIPS 140-2 on your Logger but the SmartConnectors that send events to it are not running in FIPS 140-2 mode, your deployment is not fully FIPS 140-2 compliant.

In a typical deployment, your Logger will communicate with the following components. To be fully FIPS compliant, all of these components should be FIPS enabled:

- SmartConnectors that send events to it

FIPS mode is supported on SmartConnectors running version 4.7.5.5372 and later. Follow instructions in [“Installing or Updating a SmartConnector to be FIPS compliant” on page 426](#) to ensure that your connector is FIPS compliant.

- Logger forwarders, such as ArcSight Managers to which Logger forwards events and alerts

The system to which your FIPS-compliant Logger forwards events should be FIPS compliant as well. Additionally, you need to import that system's SSL server certificate on the Logger so that Logger can communicate with it.

If you forward events and alerts to an ArcSight Manager, it needs to run ESM 4.0 SP2 or later to enable FIPS 140-2 on it. For more information, see the ArcSight ESM Installation and Configuration Guide for the ESM version you are running. Additionally, follow instructions in [“ESM Destinations” on page 291](#) to complete configuration of this setup.

- Loggers

Loggers running 45.0 or later automatically use FIPS 140-2 compliant algorithms. Therefore, no action is required on such Loggers, except enabling FIPS as described in this section.

When enabling FIPS on a software Logger, make sure that the machine on which Logger is installed is used exclusively for Logger.



Enabling FIPS 140-2 on software Logger does not make the system on which it is installed FIPS 140-2 compliant. Consult your system's documentation to determine the requirements for making the entire system FIPS 140-2 compliant.

You can enable or disable FIPS mode on Logger to suit your needs; however, you will need to restart the software Logger before the new mode will be effective.

Things to be Aware of When Enabling FIPS Mode on Logger:

- Your Logger must be set up with a CA-signed SSL certificate. For more information, see [“SSL Server Certificate” on page 417](#).
- A Logger, even when in non-FIPS mode, must use a CA-signed certificate if it is a destination of a FIPS-enabled software-based SmartConnector. Additionally, ensure that the root certificate of the CA that signed Logger's certificate is trusted on the SmartConnector. If the CA's root certificate is not trusted on the SmartConnector, follow instructions in [“Installing or Updating a SmartConnector to be FIPS compliant” on page 426](#).

- Once FIPS is enabled on your Logger, the SmartMessage receiver (if configured) stops receiving events from non-FIPS connectors if those connectors are not running version 4.7.5.5372 or later. Make sure you have the correct connectors.

To enable or disable FIPS mode:



Note

Make sure you are familiar with the configuration requirements on your Logger as described in “Things to be Aware of When Enabling FIPS Mode on Logger:” on page 425.

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **FIPS 140-2** in the Security section in the left panel.
- 3 Click **Enable** or **Disable** for the Select FIPS Mode option.
- 4 Click **Save**.
- 5 Use the following command to restart your Logger:

```
<install_dir>/current/arcsight/logger/bin/loggerd restart
```

The FIPS Status Table shows which processes and components of the Logger are FIPS enabled.

Installing or Updating a SmartConnector to be FIPS compliant

FIPS mode is supported on SmartConnectors running version 4.7.5.5372 or later.

If you are...	Then...
Installing a new SmartConnector to send events to a Logger in FIPS-compliant mode	<ol style="list-style-type: none"> 1 Download a FIPS-supported SmartConnector version from the HP Customer Support site (SSO). 2 Go to Step 1 on page 426.
Updating a SmartConnector to be FIPS compliant and the SmartConnector is not running version 4.7.5.5372 or later.	<ol style="list-style-type: none"> 1 Upgrade the SmartConnector to a FIPS-supported version. Follow instructions in the SmartConnector User's Guide to upgrade the SmartConnector. 2 Perform only Step 2a on page 427.
Updating a SmartConnector to be FIPS compliant and the SmartConnector is running version 4.7.5.5372 or later.	Perform only Step 2a on page 427 .

To make a SmartConnector FIPS compliant:

- 1 Follow device configuration steps provided in the SmartConnector's configuration guide (available from the HP Customer Support site (SSO) at <http://support.openview.hp.com>), then follow the installation procedure through installation of the core connector software (SmartConnector Installation step 2).

At Step 3 of the Connector setup, click **Cancel** to exit the setup. You must then configure the NSS DB, which is necessary for installing the connector in FIPS-compliant mode.

Once the NSS DB is configured, continue with [Step 2](#), below.

- 2 To enable FIPS Mode on the SmartConnector:

- a** Create an `agent.properties` file at the following location if it does not exist already:

```
$ARCSIGHT_HOME\current\user\agent
```

- b** Enter the following property, then save and close the file.

```
fips.enabled=true
```

3 Import Logger's Certificate on the SmartConnector:

- a** In a DOS prompt window on your SmartConnector machine, from `$ARCSIGHT_HOME\current\bin`, enter the following command to turn off FIPS mode.

```
arcsight runmodutil -fips false -dbdir  
user/agent/nssdb.client
```

- b** Export the Logger certificate file and import it to the SmartConnector's NSS DB as follows:

- i** Export Logger's certificate file from the browser you use to connect to it. Refer to your browser's Help for instructions. For example, to export a Logger's certificate file on Firefox, click **Tools > Options > Advanced > Encryption > View Certificates > Servers > *Select your Logger appliance* > Export**. Save the certificate file with a `.crt` or `.cer` extension.

- ii** Copy the certificate file you exported in the previous step (in this example, **loggercert.crt**) to the `$ARCSIGHT_HOME\current\bin` directory on the SmartConnector. From `$ARCSIGHT_HOME\current\bin`, enter the following:

```
arcsight runcertutil -A -n mykey -t "CT,C,C" -d  
user/agent/nssdb.client -i bin/loggercert.crt
```

- c** Enter the following command to re-enable FIPS mode that you turned off in Step 1:

```
arcsight runmodutil -fips true -dbdir  
user/agent/nssdb.client
```

- d** Ensure that the SmartConnector can resolve the name specified in the CN value of the Logger certificate's *Subject* field. If the name is not resolvable, add it to SmartConnector system's Hosts file.

- e** If you are installing a new SmartConnector, continue to the next step.

If you are updating your SmartConnector to be FIPS compliant, ensure that the connector's Logger destination host name is same as the CN value in the certificate's *Subject* field, and **exit this procedure**.

- 4** To return to the SmartConnector configuration wizard, enter the following from `$ARCSIGHT_HOME\current\bin`:

```
arcsight connectorsetup
```

- 5** When prompted whether you want to start in Wizard Mode, click **Yes**.

The Destination selection window is again displayed; return to your SmartConnector Configuration Guide, **SmartConnector Installation step 4** to continue the connector configuration.



Note

When configuring the connector, ensure that the connector's Logger destination host name is same as the CN value in the certificate's **Subject:** field.

For the remainder of the configuration process, see the Configuration Guide for the SmartConnector you selected to install. The specific configuration guide provides information about how to configure the device for event collection, specific installation parameters required during the configuration process, and a table of vendor-specific field mappings to ArcSight events.

Users/Groups

Use the **Users/Groups** sub-menu to configure users and user groups, and to set authentication options.

Authentication

Authentication Settings enable you to specify the settings and policies for user login sessions, password rules and lockouts, and external authentication options.

Sessions

The **Session** tab lets you specify the maximum number of simultaneous sessions for a single user account, and the length of time after which a user session is automatically logged out or a user account disabled. By default, a single user account can have up to 15 simultaneous active sessions, and a user account is logged out after 15 minutes of inactivity.

Authentication Settings

Sessions Local Password External Authentication

Session Settings

Max Simultaneous Logins/User: 15

Logout Inactive Session After: 0 hours 15 minutes 0 seconds

☒ Disable Inactive Account After: 0 days

Save

To change session settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 On the **Sessions** tab, update the parameters described in the following table.

Parameters	Description
Max Simultaneous Logins/User	The maximum number of simultaneous sessions allowed for a single user account. The default is 15 sessions .

Parameters	Description
Logout Inactive Session After	The length of time, in minutes, after which an inactive session is automatically ended. The default is 15 minutes . This value does not apply to the user interface pages accessed through the Monitor menu. If a user is on any of the Monitor menu pages and the session has been inactive for the specified number of minutes, the user's session remains active.
Disable Inactive Account After	The number of days after which an inactive user account is disabled. The default is 0 , meaning the account is never disabled.

- 4 Click **Save** to make the changes, or click another tab to cancel.

Local Password

The **Local Password** tab enables you to set password policies, such as the minimum and maximum number of characters and other password requirements.

Authentication Settings

Sessions Local Password External Authentication

Local Password Settings

Lockout Account

☒ Enable Account Lockout

Lockout Account After Failed Attempts

Remember Failed Attempts For hours minutes seconds

Lockout Account For hours minutes

Password Expiration

☒ Enable Password Expiration

Password Expires in days

Notify User Days Before Expiration

[Users Exempted From Password Expiration Policy](#)

Password Strength Rules

☒ Enforce Password Strength

Minimum Length characters

Maximum Length characters

Password Character Rules

Password must have a minimum of the following characters

Numeric [0-9] Uppercase [A-Z]

Special [!\$%^*...] Lowercase [a-z]

Password Must be At Least Characters Different From Old Password

☐ Include "Forgot Password" link on Login Screen

Save

To change the password settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **Local Password** tab.

Use the parameters described in the following table to customize your password settings.

Table 8-1 Authentication Settings, Local Password tab

Parameter	Description
Lockout Account (policy)	
Enable Account Lockout	Select the checkbox to enable user accounts to be locked out as defined by the following settings. By default, the policy is disabled .
Lockout Account After	Number of failed login attempts after which a user account is locked out. The default is 3 .
Remember Failed Attempts For	The length of time, in minutes, for which a failed login attempt is remembered. The default is 1 .
Lockout Account For	The length of time, in minutes, for which a locked out account cannot be unlocked. The default is 15 .
Password Expiration (policy)	
Enable Password Expiration	Select the checkbox to enable user passwords to expire as defined by the following settings. By default, the policy is disabled .
Password Expires in	Number of days after which the password expires. The default is 90 .
Notify User	Number of days before expiration to notify the user. Select this option to allow users to update their password before expiration. The default is 5 .
Users Exempted From Password Expiration Policy	Click the link to set the number of users whose password should never expire. For information on how to use this feature, see "Users Exempted From Password Expiration" on page 431.
Password Strength Rules (policy)	
Enforce Password Strength	Select the checkbox to enforce password policy as defined by the following settings. By default, the policy is disabled .
Minimum Length	Minimum number of characters that a password must contain. The default is 10 .
Maximum Length	Maximum number of characters that a password can contain. The default is 20 .
Password Character Rules	
Password character rules define additional character requirements to ensure password strength.	
Numeric	Minimum number of numeric characters (0-9) in a password. The default is 2 .
Uppercase	Minimum number of uppercase characters (A-Z) in a password. The default is 0 .
Special	Minimum number of non-digit and non-letter characters that are required in a password. The default is 2 .
Lowercase	Minimum number of lowercase characters (a-z) in a password. The default is 0 .

Table 8-1 Authentication Settings, Local Password tab (Continued)

Parameter	Description
Password Must be At Least N Characters Different From Old Password	Minimum number of characters by which the new password must differ by from the previous one. The default is 2 .
Include "Forgot Password" link on Login Screen	<p>Select the checkbox to enable users to reset their local password via a "Forgot Password" link on the login page. By default, the option is disabled.</p> <p>An SMTP server must be configured on the system, and the username must have a correct email address for this feature to work successfully.</p> <p>If an SMTP server is not set, you will not be able to reset the password because the email containing the temporary password cannot be sent.</p> <p>An email address must be specified in the user settings for the user name. The temporary password is sent to that email address. If no email address is specified or if the email address is correct, the user will not receive the email.</p> <p>For information on how to use this feature, see "Forgot Password" on page 432.</p>

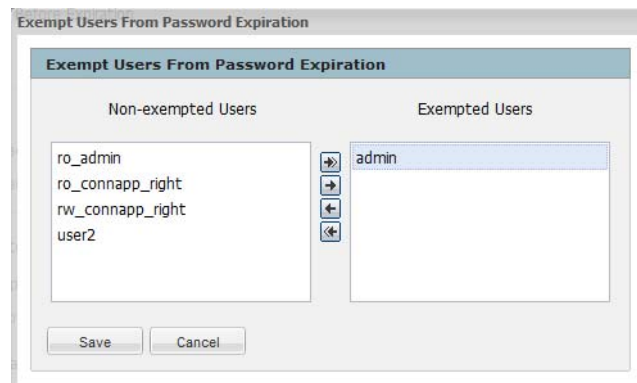
- 4 Click **Save** to save the changes, or click another tab to cancel.


Users Exempted From Password Expiration


Even though you have set a password expiration policy for most users, you may want to have a user whose password does not expire automatically.

To exempt a user from the password expiration policy:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **Local Password** tab, and then click **Users Exempted From Password Expiration Policy**.
- 4 The **Exempt Users From Password Expiration** page appears.



- 5 Select users from the **Non-exempted Users** list and click the right arrow icon  to move the selected users to the **Exempted Users** list. Do the reverse to remove users from the list of exempted users.

You can select multiple users at the same time and move them over. Or you can move all users by clicking the  icon.

- 6 Click **Save** to save the policy or **Cancel** to exit.

Forgot Password


This feature is available only if the “Include “Forgot Password” link on Login Screen” setting on the Authentication Settings page (**System Admin > Authentication > Local Password**) is set to **Yes**. By default, this setting is set to **No**. An SMTP server must be configured in order to use this feature. For more details on how to enable it, see “[Local Password](#)” on page 429.

If you forget your system password, use this feature to receive an email that provides a temporary password.

The temporary password is valid until the time specified in the email. If you do not log in within the specified time, only an administrator can reset the password to generate another temporary password.

To reset your password:

- 1 Click the **Forgot Password** link on the Login screen, as shown in the following figure.



- 2 Enter a user name on the Reset Password screen.
- 3 Click **Reset Password**.

An automated email with a temporary password is sent to the email address specified for that user.

External Authentication

Besides providing a local password authentication method, your system supports Client Certificate/CAC, LDAP, and RADIUS authentication. It is not possible to enable all authentication methods simultaneously.



Note

CAC is a form of client certificate authentication. Information on client certificate authentication applies to CAC.

From the **External Authentication** tab, use the drop-down menu to choose one of the following authentication methods:

- [Local Password](#)
- [Client Certificate Authentication](#)
- [Client Certificate and Local Password Authentication](#)
- [LDAP/AD and LDAPS Authentication](#)
- [RADIUS Authentication](#)

Local Password

This option is the default method and implements the local password policies set in the **Local Password** tab. Leave this as the default, or click **Save** if changing from another option.

Client Certificate Authentication

This authentication method requires that users authenticate using a client certificate. For each client certificate, a user account with a Distinguished Name (DN) matching the one in the client certificate must exist on your system.



All SSL client certificates used for authentication must be FIPS compliant (hashed with FIPS-compliant algorithms) even if FIPS is not enabled on your system.

To configure client certificate authentication:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **External Authentication** tab.
- 4 From the drop-down menu, choose **Client Certificate**.
- 5 **Allow Local Password Fallback** provides two options:
 - ◆ **Allow Local Password Fallback for Default Admin Only**
Select this option to allow the default admin user to log in using only a username and password if the client certificate is not available or invalid. This privilege is restricted to the default admin user only—other users must have a valid client certificate to gain access to the system. This option is enabled by default.
 - ◆ **Allow Local Password Fallback for All Users**
Select this option to allow all users to log in using their local user name and password if their client certificate is invalid or unavailable.

For more information, see [“Local Password Fallback” on page 437](#).
- 6 Click **Save**.

Client Certificate and Local Password Authentication

This authentication method requires that users authenticate using an SSL client certificate and a valid local password. *Local Password* refers to the password associated with the user credentials created in **User Management** in the **Users/Groups** section. See [“User Management” on page 438](#) for details.

A user account on your system must be defined with a Distinguished Name (DN) that matches the one in the client certificate.

For instructions on how to create a user DN, see [“Users” on page 438](#) and refer to the section called “Use Client DN” in the parameters table.



All SSL client certificates used for authentication must be FIPS compliant (hashed with FIPS-compliant algorithms) even if FIPS is not enabled on your system.

To configure client certificate and password authentication:

- 1 Click **System Admin** from the top-level menu bar.
 - 2 Click **Authentication** in the **Users/Groups** section.
 - 3 Choose the **External Authentication** tab.
 - 4 From the drop-down menu, choose **Client Certificate AND Local Password**.
 - 5 **Allow Local Password Fallback** provides two options:
 - ◆ **Allow Local Password Fallback for Default Admin Only**
This option, always enabled, allows the default admin user to log in using only a username and password.
 - ◆ **Allow Local Password Fallback for All Users**
This option is always disabled. You cannot enable it when using the **Client Certificate AND Local Password** authentication method.
- For more information, see [“Local Password Fallback” on page 437](#).
- 6 Click **Save**.

LDAP/AD and LDAPS Authentication

This authentication method authenticates users against an LDAP server. Even when LDAP is enabled, each user account must exist locally on your system. Although the user name specified locally can be different from the one specified on the LDAP server, the Distinguished Name (DN) specified for each user account must match the one in the LDAP server.



For steps on how to create a user DN, see [“Users” on page 438](#), and the parameter [“Use Client DN” on page 439](#).

To set up LDAP authentication:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **External Authentication** tab.
- 4 From the drop-down menu, choose **LDAP**.
- 5 **Allow Local Password Fallback** provides two options:

◆ **Allow Local Password Fallback for Default Admin Only**

Select this option to allow the default admin user to log in using only a username and password if LDAP authentication fails. This privilege is restricted to the default admin user only—all others must be authenticated by LDAP. This option is enabled by default.

◆ **Allow Local Password Fallback for All Users**

Select this option to allow all users to log in using their local user name and password if LDAP authentication fails.

For more information, see [“Local Password Fallback” on page 437](#).

LDAP Server has the following parameters:

Parameter	Description
Server Hostname[:port] (optional)	(Optional) Enter the host name or IP address and port of the LDAP server in the following format: <code>ldap://<hostname or IP address>:<port></code> <code>ldaps://<hostname or IP address>:<port></code> Additional steps are required for the use of LDAPS. See Using the LDAP over SSL (LDAPS) Protocol below.
Backup Server Hostname[:Port] (optional)	(Optional) Enter the backup LDAP server to use if the primary server does not respond. If the server returns an authentication failure (bad password, unknown username, etc), then the backup server is not tried. The backup server is tried only when the primary server has a communication failure. Use the same format as the primary server to specify the host name and port.
Request Timeout	The length of time, in seconds, to wait for a response from the LDAP server. The default is 10 .

6 When finished, click **Save**.

Using the LDAP over SSL (LDAPS) Protocol

When choosing the LDAPS protocol to authenticate users, make sure the following conditions are true:

- The SSL certificate for the LDAPS server has been uploaded into the trusted store.
- The external authentication method is set to “LDAP”.
- The URL for the LDAPS server(s) starts with “ldaps://”.

After uploading the SSL certificate, the **aps** process must be restarted (**System Admin > Process Status > aps Restart**).



If the aps process is not restarted, attempts to authenticate via LDAPS will fail.

RADIUS Authentication

This authentication method allows users to authenticate against a RADIUS server. Even when RADIUS authentication is enabled, each user account must exist locally on your system. The username must match the one in the RADIUS server, although the password can be different. A user must present a valid username and (RADIUS) password to be successfully authenticated.

To configure RADIUS authentication settings:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Authentication** in the **Users/Groups** section.
- 3 Choose the **External Authentication** tab.
- 4 From the drop-down menu, choose **RADIUS**.
- 5 **Allow Local Password Fallback** provides two options:
 - ◆ **Allow Local Password Fallback for Default Admin Only**
Select this option to allow the default admin user to log in using only a username and password if RADIUS authentication fails. This privilege is restricted to the admin user only—all others must be authenticated by RADIUS. This option is enabled by default.
 - ◆ **Allow Local Password Fallback for All Users**
Select this option to allow all users to log in using their local user name and password, if RADIUS authentication fails. For more information, see [“Local Password Fallback” on page 437](#).
- 6 **Update the RADIUS Server** parameters as necessary:

Parameter	Description
Server Hostname[:port]	Enter the host name and port of the RADIUS server.
Backup Server hostname[:port] (optional)	<p>(Optional) Enter the backup RADIUS server to use if the primary server does not respond. If the server returns an authentication failure (bad password, unknown username, etc), then the backup server is not tried. The backup server is tried only when the primary server has a communication failure.</p> <p>Use the same format as the primary server to specify the host name and port.</p>
Shared Authentication Secret	Enter a RADIUS passphrase.
NAS IP Address	The IP address of the Network Access Server (NAS).
Request Timeout	The length of time, in seconds, to wait for a response from the RADIUS server (in seconds). The default is 10 .
Retry Request	Number of times to retry a RADIUS request. The default is 1 .
RADIUS Protocol:	Use the drop-down menu to choose a protocol option. The default is None .

- 7 Click **Save**.

Local Password Fallback

You can use this feature to log in using your local user name and password if the external authentication (Certificate, LDAP, or RADIUS) fails, if you forgot your password to the authentication server, or if the authentication server is not available.

The Use Local Authentication allows the default admin to log in even when the remote authentication server is not available, by adding a **Use Local Authentication** checkbox to the login screen. Out-of-box, this option is enabled only for the default administrator. However, it is possible to allow local password fallback for all users. For example, you could configure the RADIUS authentication method to allow users to log in using local authentication instead of RADIUS should they fail to authenticate to the configured external RADIUS server(s).

For information on how to allow local password fallback for all users for all users, see [“Client Certificate Authentication” on page 433](#), [“LDAP/AD and LDAPS Authentication” on page 434](#), or [“RADIUS Authentication” on page 436](#).


To log in when authentication fails:

- 1 Mark the **Use Local Authentication** checkbox if the login failure was caused by failure of the external authentication.



Note

This option is only available to the default admin unless it has been enabled for other users.



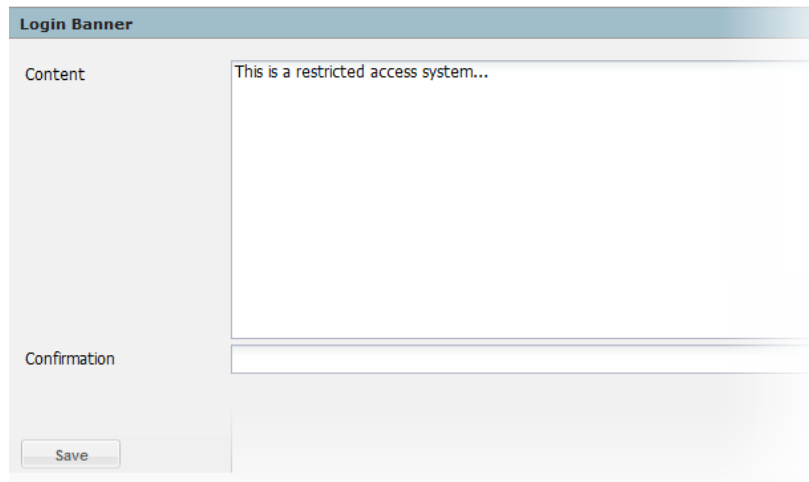
- 2 Enter your login and password and click **Login**.

Login Banner

You can customize the message on the login screen to suit your needs. The text you enter in the Content field is displayed above the Username and Password fields on the login

screen. In addition, you can enter a confirmation message that the user must click to enable the Username and Password fields.

Login Banner

The screenshot shows a web form titled "Login Banner". It has a header bar with the title. Below the header, there are two main input areas. The first is labeled "Content" and contains a text box with the text "This is a restricted access system...". The second is labeled "Confirmation" and contains an empty text box. At the bottom left of the form is a "Save" button.

You must have the "Configure Login Settings" permission enabled for your user account to edit the login banner.

To customize the login banner:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Login Banner** in the **Users/Groups** section.
- 3 Enter the text you want to display as the login banner in the Content field.

You can enter only unformatted text in this field; however, you can apply standard HTML tags to display formatted text. Loading images in this field is not allowed.

- 4 (Optional) Enter text in the Confirmation field.

If you enter text in this field, the text will be accompanied by a checkbox that the user must click to enable the Username and Password fields. For example, if you enter "Are you sure?", "Do you want to proceed?", or "I agree." in this field, the user must click the checkbox in order to log in.

- 5 Click **Save**.

User Management

The **Users** and **Groups** tabs enable you to manage users and user groups on your system. User groups are a way to enforce access control to various sections of your system.

Users

Open the **Users** tab to manage the users that can log in to your system. You can add a new user, edit user information, or delete a user at any time. You must have the appropriate System Admin group rights to perform these functions.

To add a new user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.

- 3 In the **Users** tab, click **Add** from the top left side of the page.
- 4 Enter the following parameters.

Parameter	Description
<i>Credentials</i>	
Login	The user's login name.
Password	The user's password.
Confirm Password	Reenter the users' password.
<i>Contact Information</i>	
Use Client DN	<p>If you enabled SSL client certificate or LDAP authentication, click this link to enter user's the Distinguished Name (Certificate Subject) information. The Distinguished Name should be similar to this format:</p> <p>CN=UserA,OU=Engg Team,O=ArcSight\, Inc.,L=Cupertino,C=US,ST=California</p> <p>To determine the DN, use this URL to display the certificate: <a href="https://<hostname or IP address>/platform-service/DisplayCertificate">https://<hostname or IP address>/platform-service/DisplayCertificate OR Obtain the DN information for a user from the browser that the user will open to connect to the system. For example, on Firefox, click Tools > Options > Advanced > Encryption > View Certificates > Your Certificates > Select the certificate > View.</p>
First Name	The user's first name.
Last Name	The user's last name.
Email	The user's email address.
Phone Number	(Optional) The user's phone number.
Title	(Optional) The user's title.
Department	(Optional) The user's department.
Fax	(Optional) The user's fax number.
Alternate Number	(Optional) The user's alternate phone number.
<i>Assign to Groups</i>	Select the groups to which this user belongs. This setting controls the privileges a user has on this Logger.
Notes	(Optional) Other information about the user.

- 5 Click **Save and Close**.

To edit a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 In the **Users** tab, select the user (or users) you want to edit.

- 4 Click **Edit** from the top left side of the page.
- 5 Update the user information as necessary.
- 6 Click **Save User**.

To delete a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 In the **Users** tab, select the user (or users) you want to delete.
- 4 Click **Delete** from the top left side of the page.

Reset Password

The Reset Password feature enables you to reset a user's password without knowing their password. If you are using an SMTP-configured server and have permissions to create and update users, you can reset a user's password by clicking the **Reset Password** button. An automated email is sent to the user with the new password string.

An SMTP server must be configured for the automated email containing the temporary password to be sent. If an SMTP server is not configured, the password will not be reset because an email cannot be sent.

To reset a user's password:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 In the **Users** tab, select the user (or users) whose passwords you want to reset.
- 4 Click **Reset Password** from the top left side of the page.

The user must use the temporary string to log in within the time specified in the email. If the user does not log in within the specified time, the account becomes deactivated. If the account has been deactivated, the admin must re-activate it before resetting the password.

To activate a user:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 In the **Users** tab, select the user (or users) that you want to activate.
- 4 Choose **Edit**.
- 5 Check the **Active** box.
- 6 **Save** the changes.

Groups

User groups define privileges to specific functions on your system and serve to enforce access control to these functions. For example, if you want User A to be able to run searches but not reports, assign that user to the Search group but not to the Reports group.

User groups are organized by the following types: System Admin, Logger Rights, Logger Search, Logger Reports, and Connector Appliance Rights (only on Logger platforms with integrated Connector Appliance). Each type has a pre-defined, default user group in which

all privileges for the type are enabled. To authorize a subset of the privileges for a specific group type, create a new user group and enable only the privileges you want to provide for that group. Then, assign restricted users to the newly created group.

System Admin Groups

The System Admin Group controls the system administration operations for your system, such as configuring network information, setting storage mounts, installing SSL certificates, and user management.

Read Only System Admin Group

In addition to the default System Admin Group that enables all rights (privileges), a Read Only System Admin Group is available on your system. Users assigned to this group can view System Admin settings, but cannot change them.

Refer to your system's user interface for a complete list of privileges available to this group.

Logger Rights Group

The Logger Rights Group controls the Logger application operations for your system, such as viewing the Logger dashboards and configuring all the settings in the Configuration menu (including event archives, storage groups, alerts, filters, and scheduling tasks.)

Refer to your system's user interface for a complete list of privileges available to this group.

Logger Search

The Logger Search Group controls local and peer searches through the following privileges:

- Search for events
- Search for events on remote peers

If the group is configured to allow users to run local and peer searches, users assigned to this group can perform those operations. Conversely, if the group is configured to prevent users from running local and peer searches, users assigned to this group cannot perform those operations.

Logger Reports


The Logger Reports group controls all report operations on Logger such as run, edit, delete, schedule, and view published reports.

Refer to your system's user interface for a complete list of privileges available to this group.

Managing a User Group

To create a new user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 Click the **Groups** tab.
- 4 Click **Add** from the top left side of the page.
- 5 Define the new group:
 - a In the **Group Name** field, provide a name for the group.
 - b In the **Description** field, provide a description for the group.

- c From the Group Type drop-down box, select the group type.
 - d Click the down arrow icon () next to the group type name to view and select privileges that you want to assign to the users in this group.
- 6 Click **Save and Close** to save the settings of the group, or click **Save and Edit Membership** to add users to this group.

To edit a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 Click the **Groups** tab.
- 4 Select the group that you want to edit, and click **Edit** at the top left side of the page.
- 5 Update the user group information.

If you need to edit the group's membership:

- a Click **Save and Edit Membership** to display the Edit Group Membership page.
 - b Click **Add** from the top left of the Edit Group Membership page.
 - c Select users you want to add. By default, you can add only users who do not belong to other groups of the type that you are editing. To add such users, click **Show users that belong to other <group_type> groups**.
- When you add a user who belongs to another group of the same group type as the one you are updating, that user is automatically removed from the previous group.
- d Click **OK**.
 - e Click **Back to Group List**.
- 6 Click **Save and Close**.

To delete a user group:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **User Management** in the **Users/Groups** section in the left panel.
- 3 Click the **Groups** tab.
- 4 Select the group (or groups) that you want to delete.
- 5 Click **Delete** at the top left side of the page.

Change Password

You can use the **Change Password** menu to change your password. Passwords are subject to the password policy specified by the Admin user.

To change your password:

- 1 Click **System Admin** from the top-level menu bar.
- 2 Click **Change Password** in the **Users/Groups** section in the left panel to display the Change Password for <User Name> page.
- 3 Enter the Old Password, the New Password, and enter the New Password a second time to confirm.

- 4 Click **Change Password**.

Other System Administration Information

This section contains information related to system administration that you will need to fully administer your Logger, including system health events and SNMP polling.

Monitoring System Health

You can monitor your Logger's health in these ways:

- By using a pre-defined system filter, as listed in ["System Filters/Predefined Filters" on page 128](#). The pre-defined system health filters are based on the system health events listed in ["System Health Events" on page 443](#).
- By searching for system health events in Logger's Internal Storage Group, as listed in ["System Health Events" on page 443](#). If a pre-defined system health filter does not suit your needs, you can create alerts based on the system health events.

To set up notification of system health events:

- 1 Configure the Logger's SMTP settings (see ["SMTP" on page 414](#)) or create an SNMP Destination (see ["Sending Notifications to SNMP Destinations" on page 307](#)) or Syslog Destination (see ["Sending Notifications to Syslog Destinations" on page 308](#)).
- 2 Create an Alert that uses one or more System Alert Filters or define a query that searches for the system health events in Logger's Internal Storage Group, and specify match count and threshold (see ["Alerts" on page 296](#)).
- 3 Enable the new Alert.

System Health Events

The following table lists the system health events that Logger generates. These events are also referred as Logger Internal Events because they are stored in Logger's Internal Storage Group. See [Appendix F, Examples of System Health Events on page 151](#) for examples of these events.

The pre-defined System Filters that provide system health status are based on some of these events. If a pre-defined filter does not suit your needs, create an alert using one of these events.

Starting with Logger 5.1, the format in which system health events are generated was changed to provide more meaningful information. These changes include:

- Addition of new events (for example, Current and Voltage).
- Instead of referring to all system health events as Logger Internal Event in the `name` field, meaningful names are used (for example, Fan OK, Temperature OK).
- Three severity levels for each event have been added to the `agentSeverity` field—1 (OK), 5 (Degraded), and 8 (Severe).
- The `deviceCustomString` and `deviceCustomStringLabel` field mappings have changed. Refer to a specific event to see the changes.
- Device Event Class ID (`deviceEventClassId`) and Device Event Category (`deviceEventCategory`) of the events have changed. An updated list is available in the following table.
- All hardware-related events are classified as `hardware:nnn` events, where `nnn` is a three-digit number that identifies the hardware component (for example, `hardware:13x` identifies the fan events.)

If you are upgrading from Logger 5.0 Patch 2 or earlier, any existing filters or queries based on earlier events will not work on the events collected after the upgrade. Those filters and queries will continue to work on the events collected before the upgrade. The pre-defined System Filters are compatible with both the new and the old formats.

Group	Device Event Category	Device Event Class ID
System Health Events for appliance and software Loggers		
CPU	/Monitor/CPU/Usage	cpu: 100
Disk	/Monitor/Disk/Read	disk: 102
	/Monitor/Disk/Write	disk: 103
EPS	/Monitor/Receiver/EPS/All	eps: 100
	/Monitor/Receiver/EPS/Individual	eps: 102
	/Monitor/Forwarder/EPS/All	eps: 101
	/Monitor/Forwarder/EPS/Individual	eps: 103
Memory	/Monitor/Memory/Usage/Platform	memory: 100
Network	/Monitor/Network/Usage/In	network: 100
	/Monitor/Network/Usage/Out	network: 101
Search	/Monitor/Search/Performed	search: 100
Storage Group	/Monitor/StorageGroup/Space/Used	storagegroup: 100
	Note: The size of the storage group, indicated by the "fsize" field is in GB.	

Chapter 9

Managing Connectors

This chapter applies to Logger L3x00 appliance models only.

The following topics are discussed here.

[“Connector Overview” on page 445](#)
[“Navigating the Manage Connectors Tab” on page 446](#)
[“Locations” on page 447](#)
[“Hosts” on page 451](#)
[“Containers” on page 457](#)
[“Connectors” on page 474](#)
[“Configuration Suggestions for Connector Types” on page 504](#)

Connector Overview

You can manage the configuration of these kinds of connectors:

- **Local (on-board) connectors:** Pre-installed connectors on the Logger appliance running Connector Manager. Connector Manager (software edition) ships with no pre-installed connectors.
- **Remote Connector Appliance connectors:** Pre-installed connectors on a remotely-managed Connector Appliance.
- **Software-based connectors:** Software-based connectors installed manually on a remote host.

A connector configuration consists of properties such as name and type, and a set of *parameters* that customize how the connector works in a specific environment. Parameters vary based on the type of connector; for example, a connector for a firewall has different parameters than a connector that reads an intrusion detection system database.

You can manage connectors of many types, including syslog, Simple Network Management Protocol (SNMP), BlueCoat SmartConnector via FTP, specific Intrusion Detection Systems (IDS), log files, vulnerability scanners, and operating system-specific security events. You can view the list of supported types in the drop-down menu when you configure a new connector.



The connectors you manage are configured automatically to run as *services* or *daemons*.

Individual software-based connectors are described in ArcSight documents specific to those connectors, including the connector-specific configuration guides available with each connector. You can also find general connector information in the SmartConnector User's Guide. All of these documents are available from the Protect 724 Community (<https://protect724.arcsight.com>).

Navigating the Manage Connectors Tab

The Manage Connectors tab enables you to configure and organize connectors. This section describes the user interface elements and explains how to use them effectively.

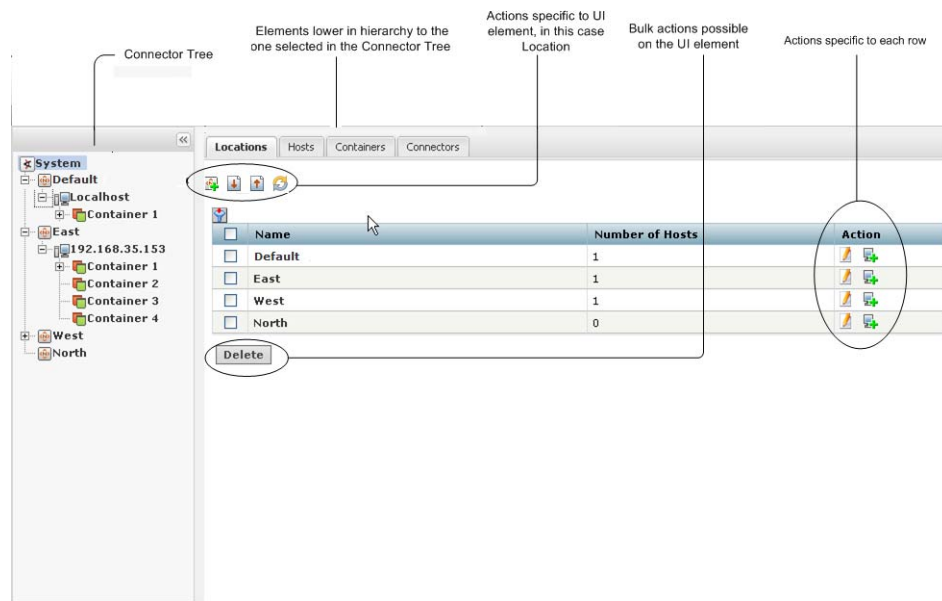



Figure 9-1 Managing Connectors

The Connector tree (the left panel of the window shown in [Figure 9-1](#)) organizes connectors into a hierarchy as follows:




Each connector you manage belongs to a container; each container belongs to a host; each host belongs to a location; and, all locations belong to root of the System.

When you click on an upper-level user interface element in the left panel, the interface displays elements lower in the hierarchy to it on the right panel. You can also perform management operations on the elements displayed on the right side.


For example, **System** provides the root (top-level) view. When you click System, all configured locations are listed in the left panel, as well as under the **Locations** tab in the right panel. You can perform various management tasks, such as editing, deleting, or adding a host, on those locations. In addition, all hosts, containers, and connectors on this system are displayed in specific tabs in the right panel. Click the **Hosts** tab to view all hosts

on the system, and click **Containers** and **Connectors** to view the respective elements and perform management operations on them. Similarly, if you select a host (from the left panel), all containers and connectors configured on that host are displayed on the right panel, as shown in the following figure.





When a container is down or a host is unreachable, the system waits for it to come online. There might be a delay of several minutes before the connector tree (in the left panel) and the Container tab (in the right panel) display.

On any user interface, you can perform three kinds of operations:

- A global operation—Listed on top of a user interface page; for example, you can upload a CSV file of locations.
- A localized operation—An operation on a single element displayed on the user interface page; for example, you can add a connector to a container by clicking the  icon in the Action column in the container's row.
- A bulk operation—A single operation performed on multiple elements on the user interface page; for example, you can upgrade multiple containers by selecting the containers (click the box to the left of the container to select it) and clicking Upgrade at the bottom of the page.



- The  icon refreshes a UI screen. This icon is available on the UI pages when relevant.
- Click the column filter icon () to display drop down lists of values on which to filter each table column. Click the checkbox in the table header to check or uncheck all checkboxes in a single column.
- When processing user provided data, Connector Appliance wizards “escape” some HTML-specific characters. Any other entered characters are not “escaped” (or validated) and are used as entered.

Locations

Location is a logical grouping of hosts. The grouping can be based on any suitable abstraction—geographical, organizational, and so on. For example, you can group all hosts in New York separately from hosts in San Francisco and label them as such. Similarly, you can group a few machines under Sales and others under Marketing.

A location can contain **any number** of hosts. **Default** location is provided on a new Connector Appliance or on a Logger appliance running Connector Manager. **Default** location is not provided on the Connector Manager (software edition).



ArcSight recommends that you do not delete the location **Default**.

You can view all locations on the system and view hosts, containers, and connectors in a location. You can add, edit, and delete a location. You can also add hosts to a location. All these procedures are described below.

Viewing All Locations

You can see all the locations that exist on the system.

To view all locations:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel.

All existing locations display on the Locations tab in the right panel.

Viewing Hosts, Containers, and Connectors in a Location

You can see all the hosts, containers, and connectors that exist in a location.

To view hosts, containers, and connectors in a location:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click the location (listed under System) from the left panel.

The hosts, containers, and connectors in the location display in the right panel, under specific tabs, as shown below.



Adding a Location


Before adding hosts, you need to add a location, which is a logical grouping of hosts.



You can also add locations in bulk using a comma-separated values (CSV) file. For more information see, [Adding Locations and Hosts from a File](#), below.

To add a location:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel.

- 3 Click  (on top of the page) in the right panel.
- 4 Enter the name of the new location and click **Next**.
- 5 Click **Done**.

Exporting and Importing Remote Management Configuration

You can create a backup of the complete remote management configuration settings on the Connector Appliance (all remote software connectors and remote Connector Appliances that are managed by the appliance) and import the configuration on another appliance.

The remote management configuration is saved in AUP format in the Remote Management AUP repository so you can download it to your local computer.


You cannot manage the same connectors using two appliances at the same time. Before importing the remote management configuration to another Connector Appliance, you need to shut down the appliance from which you exported the configuration.



Note

You can import the remote management configuration only on the same appliance model as the one from which the configuration is exported.

To export the remote management configuration:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel.
- 3 Click  (on top of the page) in the right panel.
- 4 Follow the instructions in the **Export Remotely Managed Host Wizard** to export the configuration. The remote host configuration is saved in AUP format in the Remote Management AUP repository.

After you export the remote management configuration, you need to download it to your local computer from the Remote Management AUP repository.


After you have exported the remote management configuration and have downloaded it to your local computer, you can import the configuration to another appliance.



Note

Importing the remote management configuration overwrites the current remote management configuration on the appliance.

To import the remote management configuration:

- 1 On the appliance where you want to copy the remote management configuration, click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel.
- 3 Click  (on top of the page) in the right panel.

- 4 Follow the instructions in the wizard. When selecting the type of upload, choose **Full remote management (AUP format)**.



Note

If there are no valid CA certificates for any connectors in the configuration, you see a question mark (?) next to the container that contains the connectors in the left panel. Refer to [“Resolving Invalid Certificate Errors” on page 469](#).

Adding Locations and Hosts from a File

To add hosts (and consequently, containers and connectors) in bulk, you can use a comma-separated values (CSV) file. When you add a host, the containers (and connectors) on the system are scanned automatically and the CA certificates from the containers that reside on the host are retrieved. You can manage the containers on the hosts only if it can authenticate using the certificates and the credentials. When the certificates are retrieved, you are prompted to import them.



Note


A host is **not** added if:

- Any containers on the host are down.
- If you choose not to import the certificates that are retrieved.
- Authentication fails on any of the containers.

The CSV file needs to be in the format shown in the following example. Also, ensure that an end-of-line character is included in the last line of the CSV file if the file was created on a Windows system. However, an end-of-line character is not required if the file was created on a Linux system.

	A	B	C	D	E	F
1	Location	Hostname	Port	Type	User	Password
2	East	ernie.company.com	9006	8 Containers	admin	password
3	West	elmo.company.com	9008	Software	admin	password
4						

To add locations and hosts from a CSV file:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel.
- 3 Click  (on top of the page) in the right panel to open the wizard.
- 4 Select **Remote hosts (CSV format)** and click **Next**. Follow the instructions in the wizard to upload the file.
- 5 Connector certificates are retrieved automatically so that the system can communicate with each connector in a container. The Upload CSV wizard lists the certificates. (To see certificate details, hover your mouse over the certificate.)
 - ◆ Select **Import the certificates to Connector Appliance from the containers**, then click **Next** to import the certificates and continue.
 - ◆ Select **Do not import the certificates to Connector Appliance from the containers** and click **Next** if you do not want to import the certificates. The Upload CSV wizard does not complete the upload CSV process.





Note

The Upload CSV wizard does not complete the upload if certificate download failed for any of the connectors in a container or if any of the certificates failed to import into the trust store on the system.

Editing a Location

You can edit the name of a location from the System-level page or from a specific Location page.

To edit a location:

- 1 Click **Configuration > Manage Connectors**.
- 2 From the System-level page:
Click **System** (left panel) > **Locations** tab (right panel) >  in the Action column.
From a specific Location page:
Click **System** (left panel) > *Location* >  (on top of the page, in the right panel).
- 3 Enter the new name of the location and click **Next**.
- 4 Click **Done**.

Deleting a Location

When you delete a location, the hosts, containers, and connectors that it contains are also deleted.

To delete a location:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel.
- 3 Select the location you want to delete. You can select multiple locations.
- 4 Click **Delete** at the bottom of the page, in the right panel.

Adding Hosts to a Location

See [“Adding a Host” on page 452](#).

Hosts

A host is a computer on a network, associated with an IP address, on which connectors are installed. A host can be of three types:

- The Localhost (the local Connector Appliance or the Logger appliance running Connector Manager). By default, **Localhost** exists on a brand new Connector Appliance or Logger appliance running Connector Manager; it contains a default number of containers, which are empty.
Connector Manager (software edition) does not ship with Localhost.
- A remotely-managed Connector Appliance.
- A Software-type host (a Windows, Linux, or UNIX system running software-based connectors from ArcSight). A software-type host can contain up to 20 containers.

You can view all hosts on the system, and view containers and connectors in a host. You can add, scan, delete, and edit a host. You can move a host to a different location and upgrade a host remotely. You can also add a container to a host. All these procedures are described below.

Viewing All Hosts

You can see all the hosts you are managing.

To view all hosts:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left pane. All hosts display on the Hosts tab in the right panel.

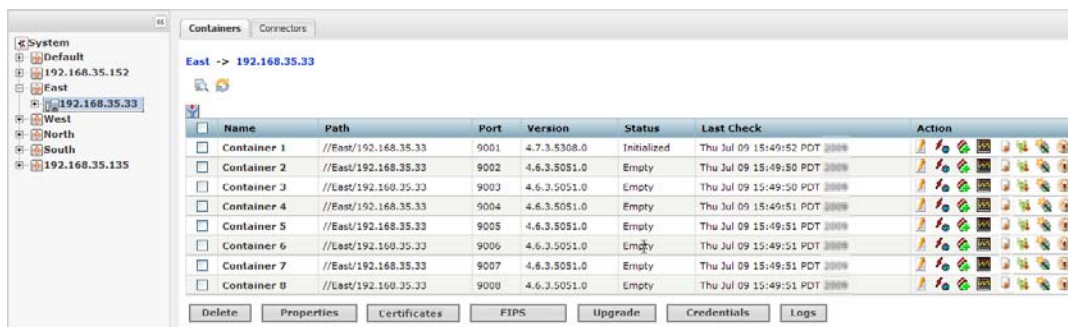
Viewing Containers and Connectors in a Host

You can see all the containers and connectors that exist on a host.

To view containers and connectors on a host:

- 1 Click **Configuration > Manage Connectors**.
- 2 In the left panel, click the location (under System) in which the host exists.
- 3 In the left panel, click the host to view the containers and connectors.

All containers display on the Containers tab and all connectors display on the Connectors tab in the right panel.



Adding a Host

By default, a local host **Localhost** exists on your Connector Appliance or Logger appliance running Connector Manager. However, Connector Appliance can manage connectors installed on other Connector Appliances and other systems such as Windows, UNIX, or Linux. To manage remote connectors, you need to add the hosts on which those connectors are running.



Connector Manager (software edition) does not ship with the default Localhost. You need to add the hosts that contain the connectors you want to manage.

When you add a host, the system also attempts to retrieve the CA certificates from the containers that reside on the host. Containers on the remote host can be managed only if

the system can authenticate using the certificates and the credentials. When the certificates are retrieved, you are prompted to import them.



Note

A host is **not** added if:

- Any containers on the host are down.
- If you choose not to import the certificates that are retrieved.
- Authentication fails on any of the containers.

You can add hosts from the System-level page or from a specific Location page.





Note

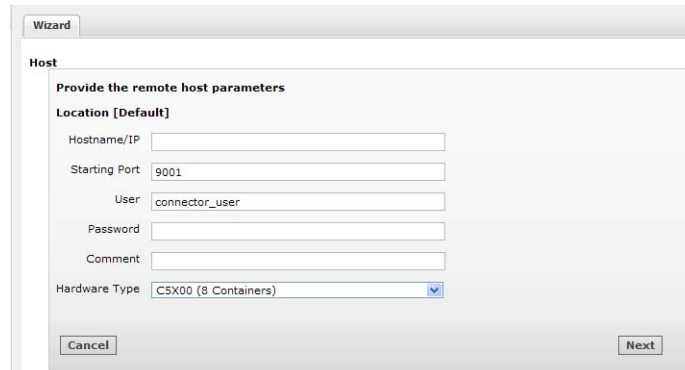
You can also add locations and hosts using a comma-separated values (CSV) file. For more information see, [“Adding Locations and Hosts from a File” on page 450](#).

When you add a remote software-type host, it is scanned automatically for the currently-running containers and the connectors associated with them. If additional containers are added to the remote host after it has been added to the system, you need to scan the host manually to detect the new containers. For information about scanning hosts, see [“Scanning a Host” on page 454](#).

To add a host:

- 1 Click **Configuration > Manage Connectors**.
- 2 From the System-level page, click **System** (left panel) > **Locations** tab (right panel) >  in the Action column.

From a specific Location page, click **System** (left panel) > *Location* (under which the host exists) >  (on top of the page, in the right panel).
- 3 On the Host Wizard form, shown below, enter values for the parameters listed in the following table and then click **Next**.



The image shows a 'Wizard' window titled 'Host'. Inside, there's a section 'Provide the remote host parameters' with a sub-section 'Location [Default]'. Below this are several input fields: 'Hostname/IP', 'Starting Port' (with '9001' entered), 'User' (with 'connector_user' entered), 'Password', 'Comment', and 'Hardware Type' (a dropdown menu showing 'CSX00 (8 Containers)'). At the bottom are 'Cancel' and 'Next' buttons.

Parameter	Description
Hostname	The hostname or IP address of the actual host.
Starting Port	Each container on a host listens on a port. Specify the starting port number. Subsequent containers will use subsequent ports.
User	The user name that the system uses to connect to the host.

Parameter	Description
Ending Port	By default, Connector Appliance scans port 9001 to port 9020 when adding a host. If you select software in the Hardware Type field, you can specify the ending port number (for example, 9003) to speed up the add host process.
Password	The password for the user name you specify.
Comment	A meaningful description for the host you are adding.
Hardware Type	<ul style="list-style-type: none">If you want to manage connectors that reside on a remote Connector Appliance, select the number of containers on that host. A host can have up to 8 containers. For the number of connectors applicable to each model type and container specifics, see the ArcSight Appliance Specifications document. This document is available at the Protect 724 Community at https://protect724.arcsight.com.If you want to remotely manage connectors running on a Windows, UNIX, or Linux system, select Software. The system can detect the presence of software-based connectors on remote hosts using the Starting Port value you specified earlier. The system scans up to 20 configurable ports from the starting port to find the "listening" connectors. Any found connectors are added into the host. For more information, see "Scanning a Host" on page 454.

- 4 Connector certificates are retrieved automatically so that the system can communicate with each connector in a container. The Add Host wizard lists the certificates. (To see certificate details, hover your mouse over the certificate.)

- ◆ Select **Import the certificates to Connector Appliance from the containers**, then click **Next** to import the certificates and add the host.
- ◆ Select **Do not import the certificates to Connector Appliance from the containers** and click **Next** if you do not want to import the certificates. Connector Appliance does not add the host.

**Note**

The Add Host wizard does not add the host if the certificate download failed for any of the connectors in a container or if any of the certificates failed to import into the trust store.

Scanning a Host

Scanning a host enables the system to detect new or removed containers from a remote **software-type** host. When a software-type host is added for the first time, it is scanned automatically for containers running at that time; however, to keep this information up-to-date, you need to scan the host manually whenever you add connectors to the remote host.

You can scan a host from the System-level page, a specific Location page, or a specific Host page.



- You can scan only software-type hosts. See [“Hosts” on page 451](#) for information about software-type hosts.
- The connectors on a software-type host need to be configured for remote management.
- A maximum of 20 connectors are scanned on port 9001 through 9020.


When you scan a host, the CA certificates from the containers that reside on the host are retrieved. The containers on the remote host can be managed only if the system can authenticate using the certificates and the credentials. When the certificates are retrieved, you are prompted to import them.



A host cannot be scanned (the scan fails) if:

- Any containers on the host are down.
- If you choose *not* to import the certificates that are retrieved.
- Authentication fails on any of the containers.

To scan a host:

- 1 Click **Configuration > Manage Connectors**.
- 2 From the System-level page, click **System** (left panel) > **Locations** tab (right panel).
From a specific Location page, click **System** (left panel) > *Location* (under which the host exists).
From a specific Host page, click **System** (left panel) > *Location* (under which the host exists) > *Host*.
- 3 Click  in the Action column for the host that you want to scan.
- 4 Click **Next** in the Host Scan wizard.
- 5 Enter values for the parameters in the following table, then click **Next**.

Parameter	Description
Starting Port	The port number on the host on which Connector Appliance starts scanning for containers.
Ending Port	The port number on the host on which Connector Appliance ends scanning for containers.
User	The user name that the system uses to authenticate with the host.
Password	The password for the user name you provide.

- 6 Connector certificates are retrieved automatically so that the system can communicate with each connector in a container. The Host Scan wizard lists the certificates. (To see certificate details, hover your mouse over the certificate.)
 - ◆ Select **Import the certificates to Connector Appliance from the containers**, then click **Next** to import the certificates and continue.

- ◆ Select **Do not import the certificates to Connector Appliance from the containers** and click **Next** if you do not want to import the certificates. The Host Scan wizard does not continue the scan.



The scan is not completed if the certificate download failed for any of the connectors in a container or if any of the certificates failed to import into the trust store.

Deleting a Host

When you delete a host, the containers and connectors that it contains are also deleted from the system that is managing the host. You can delete a host from the System-level page or from a specific Location page.

To delete a host:

- 1 Click **Configuration > Manage Connectors**.
- 2 From the System-level page, click **System** (left panel) > **Hosts** tab (right panel).
From a specific Location page, click **System** (left panel) > *Location* (under which the host exists).
- 3 Select the host you want to delete. You can select multiple hosts.
- 4 Click **Delete** on the bottom of the page.

Moving a Host to a Different Location

When you move a host, the containers and connectors that it contains are also moved. You can move a host from the System-level page or from a specific Location page.

To move a host:

- 1 Click **Configuration > Manage Connectors**.
- 2 From the System-level page, click **System** (left panel) > **Hosts** tab (right panel).
From a specific Location page, click **System** (left panel) > *Location* (under which the host exists).
- 3 Select the host you want to move. You can select multiple hosts.
- 4 Click **Move** at the bottom of the page.
- 5 Follow the instructions in the Hosts Move wizard.

Editing a Host

You cannot edit a host, however, you can delete an existing host and add a new one (as described in [“Adding Hosts to a Location” on page 451](#)) or move an existing host (as described in [“Moving a Host to a Different Location” on page 456](#)).

Upgrading a Host Remotely



The following instructions only apply to upgrading a remotely-managed Connector Appliance.

You can upgrade a single remotely-managed Connector Appliance or several remotely-managed Connector Appliances at the same time (in bulk). Follow these guidelines:

- The containers of the appliance being upgraded need to be managed on the system from which you will initiate the upgrade.

Remotely upgrading a Connector Appliance is a two-step process.

To upgrade a Connector Appliance remotely:

- 1 Upload a Connector Appliance .aup upgrade file from the HP Customer Support site (SSO) to the AUP repository.

This step is only required if the version that you want to upgrade does not already exist in the repository.


- 2 Push the .aup upgrade file to the remote Connector Appliances, as follows:
 - a Click **Configuration > Manage Connectors**.
 - b From the System-level page, click **System** (left panel) > **Hosts** tab (right panel).
From a specific Location page, click **System** (left panel) > *Location* (under which the host exists).
 - c Select the host you want to upgrade. You can select multiple hosts.
 - d Click **Upgrade** at the bottom of the page.
 - e Follow the instructions in the upgrade wizard.

Adding a Container to a Host

See ["Adding a Container" on page 459](#).

Containers

A container is a single Java Virtual Machine (JVM) that can run up to four connectors. The following illustration depicts Container 1 and the connectors it runs.




A default number of containers exist on each Connector Appliance. The number depends on the Connector Appliance hardware platform. Each container is identified with a label (Container Name) and an associated port number (9001 or higher).

Connector Manager on a Logger appliance contains one default container in the default host **Localhost**. You cannot delete this container. Connector Manager (software edition) does not contain any default containers.

You can perform many operations on containers. You can view all containers on the system and view the connectors in a container. You can add, delete, and edit a container. You can update container properties and change container credentials. You can manage certificates on a container, run a command on a container, and upgrade a container to a specific connector version. You can also view and delete container logs and run the Logfu utility. All these procedures are described below.

If you see a question mark (?) next to a container in the left panel, as shown below, the connectors in the container cannot be authenticated. The CA certificates for the connectors might be no longer valid. Refer to [“Resolving Invalid Certificate Errors”](#) on page 469.



Viewing All Containers

You can see all the containers you are managing.

To view all containers:

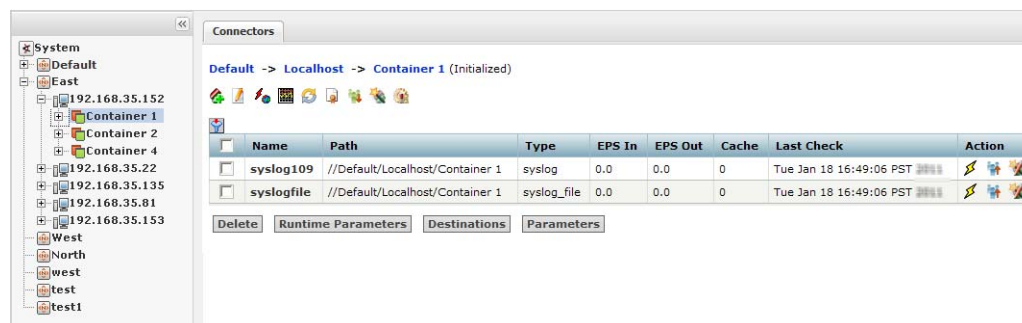
- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel. All containers display on the Containers tab in the right panel.

Viewing Connectors in a Container

You can see all the connectors in a container.

To view connectors in a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 In the left panel, click the *Location > Host* (under which the container exists) > *Container* (whose connectors you want to view). The connectors are listed on the right panel.



Adding a Container

You do not need to add a container as containers are added automatically when a new host is added to the system.

When you add a software-type host, it is scanned automatically for containers (and connectors) as described in [“Scanning a Host” on page 454](#). If you add connectors to such a host at a later date, you need to scan it manually.

Adding a Connector to a Container

See [“Adding a Connector” on page 475](#).


Editing a Container


The default names for containers are numerical, but you can change them.

To edit a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).
From the Containers page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel).

- 3 Click  in the Action column of the container whose name you want to change.

If you are on the specific Container page,  is at the top of the page.

- 4 Enter the new name in the **Name** field and click **Next**.
- 5 Click **Done**.

Deleting a Container

You can delete containers from *software-type* hosts only. All other hosts (for example, a remotely-managed Connector Appliance) have a fixed number of containers.

When you delete a container, the connectors that it contains are also deleted.

To delete a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container you want to delete. You can select multiple containers.
- 4 Click **Delete**.

Updating Container Properties

You can update existing container properties (located in the `agent.properties` file), delete them, or add new ones.

To update container properties:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container whose properties you want to update. You can select multiple containers.
- 4 Click **Properties**.
- 5 Follow the instructions in the wizard to update connector properties.

**Note**

When a property is removed, it is still visible until the container is restarted.

Changing Container Credentials

Each container has a user name and password associated with it. The default user name is `connector_user` and the default password is `change_me`. For security reasons, it is important to change these values before deploying the system in production.

To change container credentials:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container whose credentials you want to update. You can select multiple containers.
- 4 Click **Credentials**.
- 5 Follow the instructions in the wizard to update connector credentials.



Caution

This feature does not apply for containers managed by another Connector Appliance, as that appliance will not be notified of the changes. If the local system tries to communicate with the remote Connector Appliance, a credentials error occurs.

Enabling and Disabling FIPS on a Container

You can enable or disable FIPS mode on a container. When FIPS mode is enabled for a container, all the connectors in that container are in FIPS mode.

FIPS mode is supported on local, remote, and software connectors running version 4.7.5 or later. Certain connectors do not support FIPS mode. For information about which connectors do not support FIPS mode, contact customer support.



Note

Before enabling FIPS on a container that contains software connectors running as a service, review the caveats listed in document *Installing FIPS-Compliant SmartConnectors*, available from customer support.



Note


After you enable or disable FIPS mode on a container, check that the appropriate CA certificates are in the trust store of the connectors so that they can validate their configured destinations successfully. If the appropriate CA certificates are not present, you need to add them (refer to [“Managing Certificates on a Container”](#) on page 462).

To enable or disable FIPS mode on a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container on which you want to enable or disable FIPS mode. You can select multiple containers.
- 4 Click **FIPS**, then click **Next** to run the FIPS Management wizard.



- 5 Click **Enable FIPS Mode** or **Disable FIPS Mode**, then click **Next**.

If FIPS mode is already enabled or disabled on the container, the FIPS Management wizard indicates this on the Summary page.



- 6 Check that the appropriate CA certificates are in the trust store so that the connectors in the container can validate their configured destinations successfully. If necessary, add the appropriate certificates to the container. Refer to [Managing Certificates on a Container](#) below.

Managing Certificates on a Container

Connectors require a Certificate Authority (CA) issued or self-signed SSL certificate to communicate securely with a destination. The Certificate Management wizard, available

from the Containers tab, helps you add and remove certificates on a container. Using the wizard, you can:

- Enable or disable a demo certificate on a container.
You can enable a demo certificate on a container that is in non-FIPS mode only.
- Add a certificate on a container.
- Add a CA Certs file on a container.
You can add a CA Certs file on a container that is in non-FIPS mode only.
- Remove a certificate from a container.

From the Containers tab and the Connectors tab, you can view details about the certificates applied to a container. See [“Viewing Certificates on a Container” on page 468](#).

For information about resolving invalid certificates, see [“Resolving Invalid Certificate Errors” on page 469](#).


Adding CA Certificates on a Container

You can add a single CA certificate on a container that is in FIPS mode or non-FIPS mode.



Note

Whenever you enable or disable FIPS mode on a container, check that the required certificates are present in the trust store and add them if necessary.

Hover your mouse over a container name to see the type of certificate applied to it. Click the  icon to display a list of the certificates available on the container.

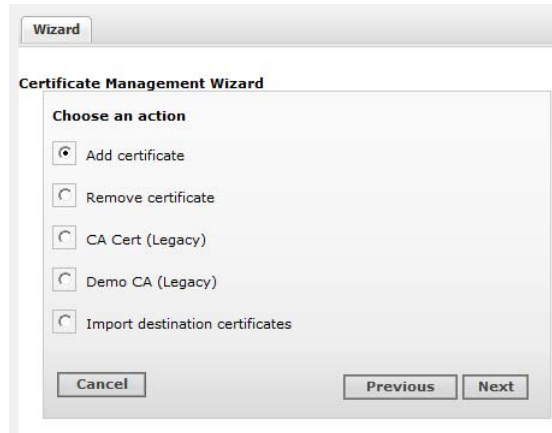
Before you follow the following procedure, make sure that the certificate you want to apply is loaded in the CA Certs repository.

To apply a single CA certificate on a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:


User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container to which you want to add the certificate. You can select multiple containers.
- 4 Click **Certificates**, then click **Next** to run the Certificate Management wizard.
- 5 Click **Add Certificate** to add a certificate.



- 6 Follow the instructions in the wizard.

If a container is down or a connector is running an older build, the wizard reports errors in the progress bar and on the Summary page.



Removing CA Certificates from a Container

You can remove CA certificates from a container when they are no longer needed. When you remove a CA certificate, the certificate is removed from the container's trust store; but it is **not** deleted from the repository.



Use caution when deleting certificates. When you delete a certificate on a container but the connector destination is still using that certificate, the connector can no longer communicate with the destination.

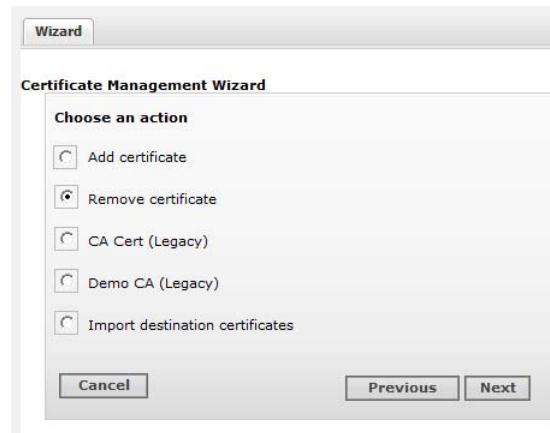
To remove CA certificates from a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container from which you want to remove the CA certificates. You can select multiple containers.

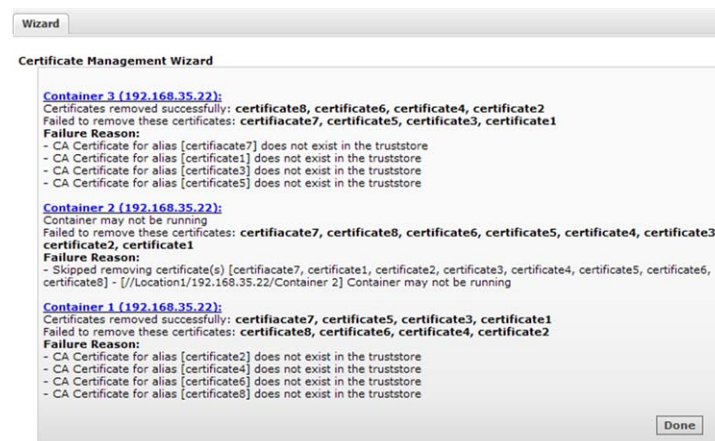
- 4 Click **Certificate**, then click **Next** to run the wizard.
- 5 Click **Remove certificate** and click **Next**.



- 6 Select one or more certificates from the certificate list and click **Next**.

The certificates are removed from the list of certificates and no longer used. When you remove a certificate from a container in FIPS mode, the container restarts automatically.

The Certificate Management wizard displays the certificates that are removed successfully in a comma-separated list. Certificates that cannot be removed are shown in a comma-separated list together with a reason why the certificate removal failed.



Adding a CA Certs File on a Container

You can add a CA Certs file on any container that is in non-FIPS mode.



Caution

When you apply a CA Certs file, the entire trust store on the container is overwritten. All previously-added certificates are overwritten.

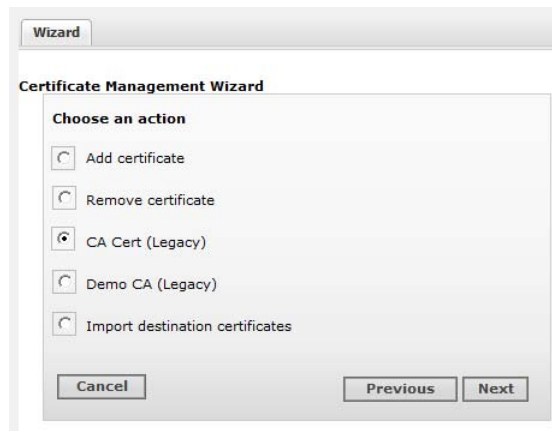
Before you follow the procedure below, make sure that the CA Certs file you want to add is loaded in the CA Certs repository.

To add a CA Certs file on a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container to which you want to add the CA Certs file. You can select multiple containers.
- 4 Click **Certificates**, then click **Next** to run the wizard.
- 5 Click **CA Cert (Legacy)**. You can add a CA Certs file to a container only if it is in non-FIPS mode.



- 6 Follow the instructions in the wizard.

After the CA Certs file has been added to a container, the container restarts automatically.

Enabling or Disabling a Demo Certificate on a Container

You can use the demo certificate on a container for testing purposes. By default, the demo certificate on a container is disabled. You can enable the demo certificate temporarily for testing purposes on a container that is non-FIPS mode.



Note

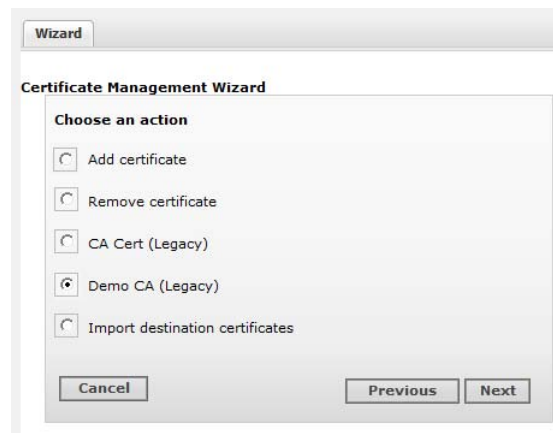
- Enable a *demo* certificate on a container in non-FIPS mode for testing purposes only. Using a demo certificate in a production environment is a serious security issue because the demo certificate is not unique.
- Hover your mouse over a container name to see the type of certificate applied to it.

To enable or disable a demo certificate on a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Select the container to which you want to apply the demo certificate. You can select multiple containers. All the containers need to be in non-FIPS mode.
- 4 Click **Certificates**, then click **Next** to run the Certificate Management wizard.
- 5 Click **Demo CA (Legacy)**, then click **Next**.



- 6 Follow the instructions in the Certificate Management wizard.


After you add the demo certificate on a container, the container restarts automatically.

Adding Multiple Destination Certificates to a Container

You can add multiple destination certificates to a container that is in FIPS mode or non-FIPS mode.



Whenever you enable or disable FIPS mode on a container, check that the required certificates are present in the trust store and add them if necessary.

Click the  icon to display a list of the certificates available on the container.

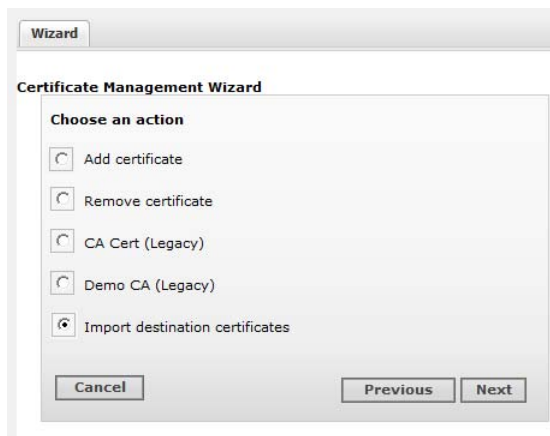
To apply multiple destination certificates to a container:

- 1 Click **Configuration > Manage Connectors**.

- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).


- 3 Select the container to which you want to add the certificates. You can select multiple containers.
- 4 Click **Certificates**, then click **Next** to run the Certificate Management wizard.
- 5 Click **Import destination certificates** to add a certificate.

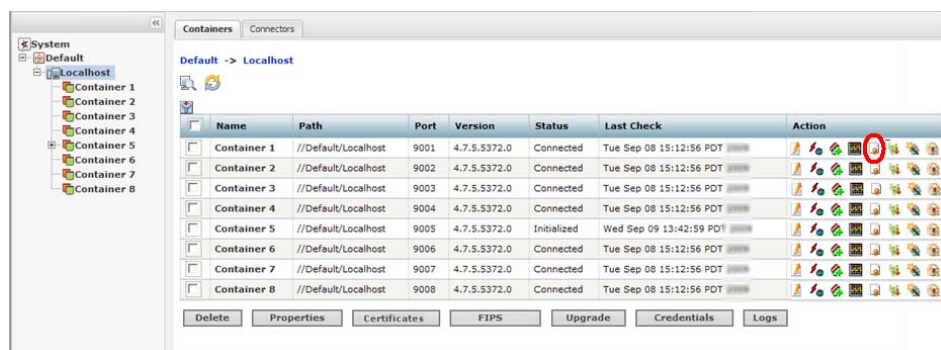



- 6 Follow the instructions in the wizard.

Viewing Certificates on a Container

From the Containers tab or the Connectors tab, you can display a list of the CA certificates applied to a container and view the details for a particular certificate in the list.

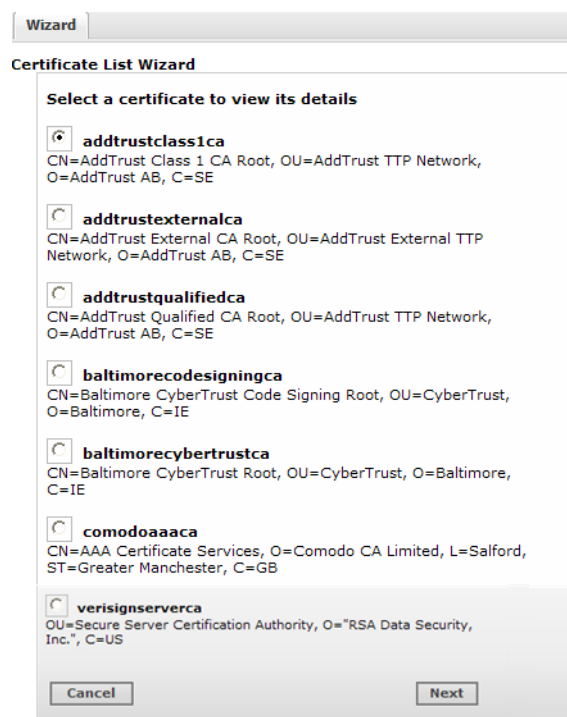
- On the **Containers** tab, click the  icon in the **Action** column for the container whose certificates you want to view.



- On the **Connectors** tab, select the  icon at the top of the page.




The Certificate List wizard displays the certificates applied to a container. To see details about a certificate, select the certificate and click **Next** at the bottom of the page.





Resolving Invalid Certificate Errors

If no valid CA certificates exist for the connectors in the container, you see a question mark (?) next to the container in the left panel, as shown below.



To resolve the invalid certificate error:

- 1 Click the container name in the left pane to view the certificate error on the Connectors tab.
- 2 Click the  icon to run the Certificate Download wizard.



- 3 Follow the instructions in the wizard to import the valid certificates.



Running a Command on a Container

You can run commands on a container to configure memory settings, pull an OPSEC certificate, or restart the container.

To run a command on a container:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).
From the Container page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Container .

- 3 Click  in the Action column of the container.
If you are on the specific Container page,  is at the top of the page.
- 4 Select the command you want to run and click **Next**.
- 5 Enter values for the parameters that the user interface displays and click **Done**.

Upgrading a Container to a Specific Connector Version

All connectors in a container are upgraded to the version you select.



Note

You can't upgrade the same container more than once within a short period of time. After you upgrade a container, wait at least 15 minutes before upgrading it again.

To upgrade a container to a specific connector version:

- 1 Upload a connector build AUP from the HP Customer Support site (SSO) to the AUP (Upgrade) repository.

This step is only required if the build does not already exist in the AUP (Upgrade) repository.

- 2 Apply the connector build to a container, as follows:

- a Click **Configuration > Manage Connectors**.
- b Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > Location (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > Location (left panel) > Host (left panel) > Containers tab (right panel).

- c Select the container that you want to upgrade. You can select multiple containers for a bulk upgrade.
- d Click **Upgrade**.
- e Select the version to which you want to upgrade the selected containers and click **Next**.



- On a slow network or when the system is under a particularly heavy load, the upgrade might be interrupted by a session timeout. To prevent this interruption, you can upload the .aup file to a higher-performance system if one is available, then push the result to the lower-performance system.
- If you are upgrading an empty container, the system creates a temporary connector during the upgrade process. You can safely ignore this temporary connector; it is deleted shortly after being created.
- Empty connectors can be upgraded from versions **5.1.2 and after**. Upgrading empty connectors is not supported in previous versions.

Viewing Container Logs

You can retrieve and view the log files for a container. The log files are in .zip format.

To view container logs:

- 1 Load the logs to the Logs repository.

If the logs that you want to view are already in the Logs repository, skip this step.

- a Click **Configuration > Manage Connectors**.

- b** Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- c** Select the container whose logs you want to view. You can select multiple containers.

- d** Click **Logs**.

The logs are loaded to the Logs repository. If you selected multiple containers, a log file for each container is loaded.

- 2** Retrieve and view the logs:

- a** Click **Configuration** > **Repositories** from the top-level menu bar.

- b** Click **Logs**.

- c** Click  to retrieve the log files (in .zip format) you want to view.

Deleting Container Logs

To delete a container log file, click **Configuration** > **Repositories** > **Logs** > from the top-level menu bar. In the right panel, click  next to the log files you want to delete.

Running Logfu on a Container

The **Logfu** utility is a diagnostic tool that parses ArcSight logs to generate an interactive visual representation of the information contained within the logs.

When event flow problems occur (with a connector or the connected device), it is useful to have a visual representation of what happened over time. You can use Logfu to analyze this behavior.



To run Logfu on a container:

- 1** Click **Configuration** > **Manage Connectors**.

- 2** Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).

User Interface Options	Path
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).
From the Container page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- Click  in the Action column of the container. A separate window is displayed. If you are on the specific Container page,  is at the top of the page.

The system proceeds to retrieve and analyze system data logs. After this process is complete, a group of panels appear in the window.

- From the **Group** box, choose which type of data you would like to view. The Group box lists all connectors within the chosen container, plus many other types of data such as memory usage, and transport rates and logs.

Choose one of the Group box **data points**. Depending on which data point you chose, a list of fields appears in the Field box below.
- Choose a **field** to view. A graphic chart appears in the Chart box, providing rate and time information. The key at the bottom of the Chart box defines the data points mapped in the chart.
- If you need to choose a different data point for analysis, click **Reset Data**.


Running Diagnostics on a Container


You can run certain diagnostics on a local or remote container. Currently, the **Edit a File** diagnostic action only is available:

To run diagnostics on a container:

- Click **Configuration > Manage Connectors**.
- Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).
From the Container page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel). The Connectors tab displays.

- To open the Container Diagnostics wizard:
 - From the **Container** tab, click  in the **Action** column.

- ◆ From the **Connectors** tab, click  at the top of the page.
- 4 Follow the steps in the wizard:
 - a Select the action you want to take on the selected container:
 - Select **Edit a configuration file** to edit a file in the `user/agent` folder on the container with the extension `.properties`, `.csv` or `.conf`.
 - Select **Edit a user file** to edit any file (except binary files, such as `.zip`, `.jar`, or `.exe`) in the `user/agent` folder on the container.
 - b From the list of available files, select the file you want to edit. The file displays in the Edit File panel. Make your edits, then click **Next** to save your edits and restart the container.



On Mozilla Firefox, if the text is underlined with red lines, right click on the text area and uncheck **Check Spelling**.



When you click **Next**, Connector Appliance saves the updated file in the `user/agent` folder on the container; the original file is overwritten.

- c Click **Done** to close the Diagnostics wizard.

Connectors

A connector (also known as a SmartConnector) is an ArcSight software component that collects events and logs from various sources on your network. A connector can be configured on a Logger appliance running Connector Manager, on a Connector Appliance, or can be installed on a computer on your network and managed remotely. For a complete list of supported connectors, go to the HP Customer Support site (SSO).

You can perform many operations on connectors. You can view all the connectors you are managing and add, remove, and edit a connector. You can update connector and table parameters, add and remove connector destinations, and edit destination parameters and runtime parameters. You can send a command to a connector or a destination, and run the Logfu utility. All these procedures are described below.



Whenever applicable, the above listed operations can be performed on more than one connector at a time. Each procedure described in this section indicates if multiple connectors can be selected when performing a procedure.

Viewing all Connectors

You can see all the connectors you are managing.

To view all connectors:

- 1 Click **Configuration > Manage Connectors**.
- 2 Click **System** in the left panel. The connectors display on the Connectors tab in the right panel.

Adding a Connector

Before you add a connector, review the following important information.

- Make sure that the container, host, and location to which you want to add the connector exist on the system. If any of these elements do not exist, first create them using procedures described in [“Adding a Location” on page 448](#), [“Adding a Host” on page 452](#), and [“Adding a Container” on page 459](#).

- Follow the configuration best practices described in [“Configuration Suggestions for Connector Types” on page 504](#).

If you are configuring the Check Point OPSEC NG Connector, see [“Configuring the Check Point OPSEC NG Connector” on page 505](#) and refer to the SmartConnector Configuration Guide for Check Point OPSEC NG.

If you are configuring a database connector that requires the MS SQL Server Driver for JDBC, follow instructions in [“Adding the MS SQL Server JDBC Driver” on page 507](#).



Caution

This connector type has special requirements concerning JDBC and authentication setup. It is important that you refer to the SmartConnector Configuration Guide for Microsoft SQL Server Multiple Instance Audit DB for this important information before installing the connector.

- If you are adding a software-based connector, make sure that the username and password for the connector match the username and password for the container to which you are adding the connector. Refer to [“Changing Container Credentials” on page 461](#).

- File-based connectors use the Network File System (NFS) or the Common Internet File System (CIFS).

For the file-based connectors on a Windows system, a CIFS share needs to be configured before you add those connectors. For information on creating a CIFS Mount, see [“Remote File Systems” on page 372](#).

For all other connectors, an NFS Mount needs to be established before the connector can be added. For information on creating an NFS Mount, see [“Remote File Systems” on page 372](#).

- For file-based FlexConnectors, make sure that an NFS Mount is established and a repository is created on the system before you add the connector. In addition, when entering the connector parameters, type the configuration file name without an extension in the Configuration File field. The extension `.sdkrfilereader.properties` is appended automatically.

To add a Connector:




Tip


If you are adding a connector for the Check Point FW-1/VPN-1 system, see a more detailed procedure in [“Configuring the Check Point OPSEC NG Connector” on page 505](#).

- 1 Click **Configuration > Manage Connectors**.

- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).
From the Container page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Click  in the Action column of the container to run the wizard to configure a connector.


If you are on the specific Container page,  is at the top of the page.

- 4 Select a connector type from the pull-down list of available types. Click **Next**.
- 5 Enter basic parameters for the connector. Parameters vary based on the connector type. You can hover the mouse pointer over a field for more information. When all fields have been entered, click **Next**.



When entering parameters that include a file path, enter the path in POSTIX format (for example, `/folder/filename`). If you enter the path in DOS/NTFS format (for example, `\folder\filename`), the backslash (`\`) is included as part of the file name and the path will be incorrect.

For file-based connectors on Windows systems, specify the name of the CIFS mount point you created for the connector, as shown in the following example. (You need to specify `/opt/mnt/CIFS_share_name`.)



Some connectors include table parameters. For example, the Microsoft Windows Event Log includes parameters for each host in the domain and one or more log types (security, application, system, directory service, DNS, file replication, and so on). You can import table parameters from a CSV file. See [“Adding Locations and Hosts from a File” on page 450](#) for the file format. You can import a CSV file that was exported from another connector as long as you export and import the CSV file from the same container. If the CSV file was exported from a different container, you need to change

the secret parameters, such as the password, which appear in obfuscated format in the CSV file to plain text before you import the CSV file.



For connectors that query Microsoft Active Directory to detect devices (for example, Microsoft Windows Event Log - Unified), if the "Network Security: LDAP Server Signing Requirements" policy is set to "Signing Required" on the Domain Controller, Connector Appliance will be unable to connect to the Active Directory or browse for devices. You see an error when selecting **Windows Host Browser** as the connector device browser type.



For detailed information about individual connector parameters, refer to the specific ArcSight SmartConnector Configuration Guide for the type of connector chosen. The configuration guide also describes how to set up the source device for use with the connector.

- 6 Choose a primary destination for the connector and enter destination-specific parameters on the following page(s), then click **Next**. Destinations can be:

- ◆ ArcSight Logger SmartMessage (encrypted)
- ◆ ArcSight Manager (encrypted)
- ◆ CEF Syslog (cleartext, that is, unencrypted)



For containers running v5.1.2.5823 and later, Connector Appliance retrieves the certificate for the destination automatically and displays the certificate summary. To see certificate details, hover your mouse over the certificate.

- Select **Import the certificate to the connector from the destination**, then click **Next** to import the certificate and continue.
- Select **Do not import the certificate to the connector from the destination** and click **Next** if you do not want to import the certificate. The destination will not be added.

For containers running v5.1.2 and earlier, upload the certificate on the container and then add the destination.

Note: FIPS Suite B mode is not supported. Connector Appliance cannot download a manager certificate in Suite B mode.

- 7 Enter connector details:

Parameter	Description
Name	A descriptive name for this connector.
Location	The location of the connector (such as the hostname).
Device Location	The location of the device that sends events to the connector.
Comment	Additional comments.



Configuring a connector can take some time; the connector might initially display *Down* while it is restarting.

- 8 Click **Done**.

Editing Connector Parameters

ArcSight supports a large number of connector types to gather security events from a variety of sources, including syslog, log files, relational databases, and proprietary devices. Accordingly, configuration parameters vary widely depending on the type of connector being configured. You can edit parameters (simple and table) for a specific connector or for multiple connectors at the same time.

Updating Simple Parameters for a Specific Connector

The following procedure describes how to update simple parameters for a specific connector. To update *table* parameters for a specific connector, see [“Updating Table Parameters for a Specific Connector” on page 479](#).

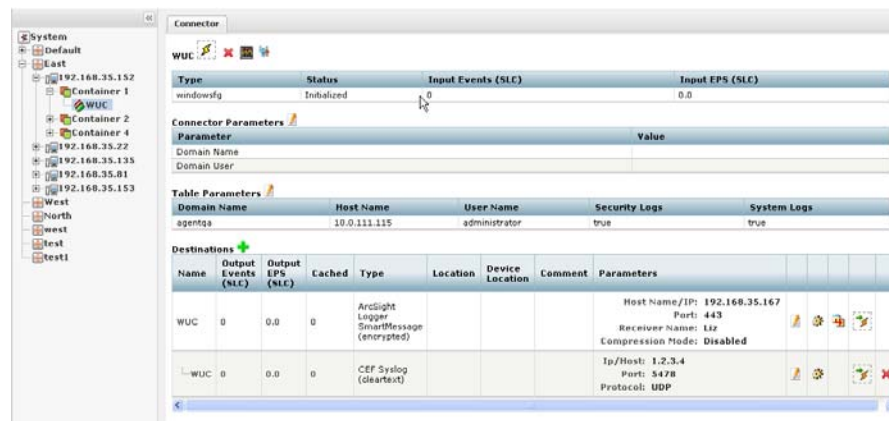
To update both simple and table parameters for multiple connectors at the same time, see [“Updating Simple and Table Parameters for Multiple Connectors” on page 480](#).

To update parameters for a specific connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click () next to the **Connector Parameters** link.



Clicking the heading **Connector Parameters** toggles between displaying and hiding the information in the Connector Parameters section.

- 4 Modify parameters as necessary and click **Next**.



Note

- Configuration parameters depend on the type of connector being configured.
- When editing parameters that include a file path, enter the path in POSTIX format (for example, `/folder/filename`). If you enter the path in DOS/NTFS format (for example, `\folder\filename`), the backslash (\) is included as part of the file name and the path will be incorrect.

- 5 Click **Done** when complete.

The updated parameters display in the Connector Parameters section of the Connector page.


Updating Table Parameters for a Specific Connector

Certain connectors, such as the Microsoft Windows Event connector, have table parameters. You can update the table parameters for a specific connector when necessary.

To update table parameters for a specific connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click () next to the **Table Parameters** link.



Note

Clicking the heading **Table Parameters** toggles between displaying and hiding the information in the Table Parameters section.

- 4 Modify parameters as necessary and then click **Next**.

- ◆ To add more rows of parameter information, click the **Add Row** link.

- ◆ You can use an Excel-compatible program to prepare a comma-separated values text file with the information and click the **Import File** button to load the entire table at once. The file needs to be in the same format as the rows shown on the Update Table Parameters page and needs to include a header row with parameter labels in the order shown on that page. For fields that require checkbox values, enter True or False as the value. An example is shown below.

	A	B	C	D	E	F
1	Domain Name	Host Name	User Name	Password	Security Logs	System Logs
2	test	1.1.1.1	admin	password	TRUE	FALSE
3	test2	1.1.1.1	admin	password	TRUE	FALSE



Note

You can import a CSV file that was exported from another connector as long as you export and import the CSV file from the same container. If the CSV file was exported from a different container, you need to change the secret parameters, such as the password, which appear in obfuscated format in the CSV file to plain text before you import the CSV file.

- ◆ To export the table parameters to a CSV file for use as a backup or to import on another Connector Appliance, click the **Export File** button.

- 5 Click **Done** when complete.

The updated table parameters display in the Table Parameters section of the Connector page.

Updating Simple and Table Parameters for Multiple Connectors

If you have multiple connectors of the same type, you can change the simple and table parameters for all the connectors at the same time.

To edit parameters for multiple connectors at once:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Select the connectors whose parameters you want to update.



Note

The connectors must be the same type; for example, you can change the parameters for several syslog connectors at the same time; however, you cannot change the parameters for several syslog and several SNMP connectors at the same time.

- 4 Click **Parameters**.
- 5 Follow the instructions in the wizard.
 - ◆ You can choose to modify the simple parameters for all the selected connectors at once or modify the simple parameters per connector.
 - ◆ If the connectors have table parameters, the table parameters are displayed so that you can modify them. If you have many table parameters to modify for multiple connectors, you can import the parameters from a CSV file (for information about adding rows and CSV file format, see [Step 3 on page 479](#)). You can also export the table parameters to a CSV file for use as a backup or to import on another Connector Appliance.



When you update parameters for connectors that are of different versions, the newer connectors might have additional parameters. In this case, only the parameters that are the same for all connectors are displayed for updating.




Managing Destinations

Connectors can forward events to more than one destination, such as ArcSight Manager and ArcSight Logger. You can assign one or more destinations per connector. You can assign multiple destinations to a connector and specify a failover (alternate) destination in the event that the primary destination fails.

The following procedures describe how to perform these actions on a specific connector or for multiple connectors at the same time:

- Add a primary or failover destination
- Edit destination parameters and destination runtime parameters
- Remove destinations
- Re-register destinations
- Manage alternate configurations for a destination
- Send a command to a destination



- You cannot configure two connectors with the same ArcSight Manager destination if the destination (connector) name and location used for configuration is the same.
- Logger receivers do not support encrypted data.
- You cannot use the **Edit** button () to change or add a connector destination. Its purpose is to change destination parameters. To add a new destination, remove the unwanted destination configuration () and create a new one ()


Adding a Primary Destination to a Specific Connector

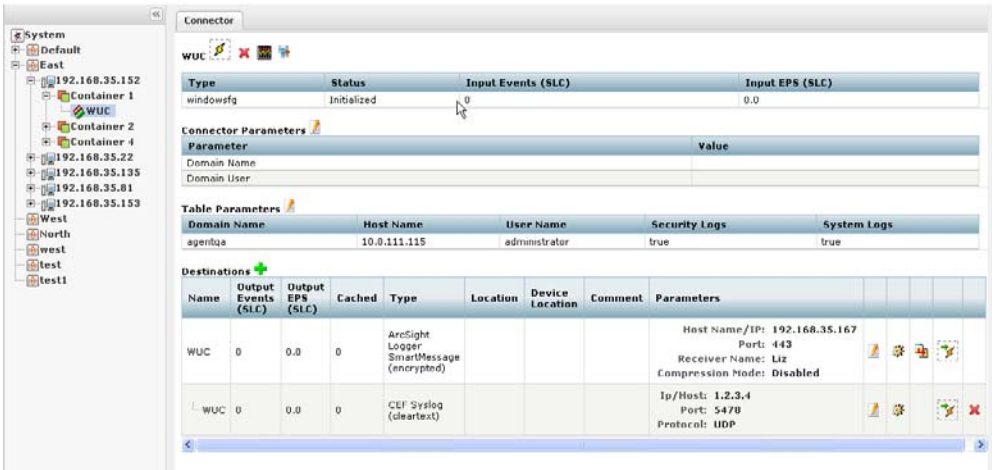
When you add a primary destination to a connector, you need to enter details for the destination, such as the destination hostname and port used.

To add a primary destination to a connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click () next to the **Destinations** link.



Note

Clicking the **Destinations** heading toggles between displaying and hiding the information in the Destinations section.

- 4 Follow the steps in the wizard.

You can either select an existing destination or add a new destination. If you are adding a new destination, select the destination type and enter parameters for the destination.



For containers running v5.1.2.5823 and later, Connector Appliance retrieves the certificate for the destination automatically and displays the certificate summary. To see certificate details, hover your mouse over the certificate.

- Select **Import the certificate to the connector from the destination**, then click **Next** to import the certificate and continue.
- Select **Do not import the certificate to the connector from the destination** and click **Next** if you do not want to import the certificate. The destination will not be added.

For containers running v5.1.2 and earlier, upload the certificate on the container and then add the destination.

Note: FIPS Suite B mode is not supported. Connector Appliance cannot download a manager certificate in Suite B mode.

- 5 Click **Done** when complete.

Adding a Failover Destination to a Specific Connector

Each destination can have a failover destination that is used if the connection with the primary destination fails.




UDP connections cannot detect transmission failure; use Raw TCP for CEF Syslog destinations.

To add a failover destination:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click () in the Destinations section to display the Add Connector Destination wizard.

- 4 Follow the steps in the wizard to select from available destinations and enter the destination details.



For containers running v5.1.2.5823 and later, Connector Appliance retrieves the certificate for the destination automatically and displays the certificate summary. To see certificate details, hover your mouse over the certificate.

- Select **Import the certificate to the connector from the destination**, then click **Next** to import the certificate and continue.
- Select **Do not import the certificate to the connector from the destination** and click **Next** if you do not want to import the certificate. The destination will not be added.

For containers running v5.1.2 and earlier, upload the certificate on the container and then add the destination.

Note: FIPS Suite B mode is not supported. Connector Appliance cannot download a manager certificate in Suite B mode.

Adding a Primary or Failover Destination to Multiple Connectors

You can add a primary or failover destination to several connectors at the same time.

To add a primary or failover destination to more than one connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Select all connectors to which you want to assign a destination.
- 4 Click **Add Destinations** at the bottom of the page to open the wizard.
- 5 Select **Add a destination** and click **Next**.
- 6 Choose between a creating a new destination or selecting an existing destination, then click **Next**.

If you choose to **create a new destination**, select the destination type and then provide the destination parameters.

If you choose to **select an existing destination**, select a destination from the list.



Connector Appliance retrieves the certificate for the destination automatically and displays the certificate summary. To see certificate details, hover your mouse over the certificate.

- Select **Import the certificate to the connector from destination**, then click **Next** to import the certificate and continue.
- Select **Do not import the certificate to the connector from the destination** and click **Next** if you do not want to import the certificate. The destination will not be added.

Note: FIPS Suite B mode is not supported. Connector Appliance cannot download a manager certificate in Suite B mode.

7 Define the destination function by choosing between a primary or failover destination.

If you choose **Primary destination**, click **Next** to update the configuration.

If you choose **Failover destination**:

- Select the primary destination that applies to your failover.
- Click the checkbox in the table header to modify all of the displayed connectors.
- Click **Next** to update the configuration.

8 Click **Done** when complete.

Removing Destinations

You can remove a destination from a connector at any time. The following procedures describe how to remove a single destination from a specific connector and how to remove multiple destinations from one or more connector.


To remove a single destination from a *specific* connector:

- Click **Configuration > Manage Connectors**.
- Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 In the Destinations section, click  for the destination you want to remove.



The  shows in the Destinations table only if more than one destination is listed.

- 4 When prompted, confirm the removal.

To remove *multiple* destinations from one or more connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Select the connectors whose destinations you want to remove.
- 4 Click the **Destinations** button to open the wizard.
- 5 Select **Remove destinations** and click **Next**.
- 6 Follow the instructions in the wizard and click **Done** when complete.

Re-Registering Destinations

At certain times, you might need to re-register the destinations for one or more connector; for example, after you upgrade ESM, or if a Logger appliance or ESM appliance becomes unresponsive.

To re-register destinations for one or more connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).

User Interface Options	Path
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Select the connectors whose destinations you want to re-register.
- 4 Click the **Destinations** button to open the wizard.
- 5 Select **Re-register destinations** and click **Next**.
- 6 Follow the instructions in the wizard and click **Done** when complete.

Editing Destination Parameters

The following procedures describe how to edit destination parameters for a specific connector and how to edit destination parameters for multiple connectors at the same time.




You cannot change the connector type; however, you can remove the unwanted connector configuration and create a new one.

To edit destination parameters for a specific connector:




- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 In the Destinations section, click () next to the destination you want to edit to display the Edit Destination Parameters page.



Caution

You cannot use the **Edit** button () to change or add a connector destination. Its purpose is to change destination parameters. To add a new destination, remove the unwanted destination () and create a new one ().

4 Make your changes and click **Next**.

5 Click **Done** when complete.

To edit destination parameters for multiple connectors:

1 Click **Configuration > Manage Connectors**.

2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

3 Select the connectors whose destination parameters you want to edit.

4 Click **Destinations** to open the wizard.

5 Select **Edit a destination** and click **Next**.

6 Follow the instructions in the wizard and click **Done** when complete.

Editing Destination Runtime Parameters



The runtime parameters for a destination enable you to specify advanced processing options such as batching, time correction, and bandwidth control. The parameters you can configure are listed in Destination Runtime Parameters appendix at the end of this guide . All the parameters listed in that table are not available for all destinations. The user interface automatically displays the parameters valid for a destination.


The following procedures describe how to edit the runtime parameters for a specific connector and how to edit the runtime parameters for multiple connectors at the same time.

To edit destination runtime parameters for a specific connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 In the Destinations section, click  next to the destination whose runtime parameters you want to edit.
- 4 Click  next to the alternate configuration that you want to edit.

If you have not set up alternate configurations, click  next to the **Default**. For more information about alternate configurations, see [“Managing Alternate Configurations” on page 490](#).

- 5 Specify or update values for the listed parameters and click **Save**.

To edit destination runtime parameters for multiple connectors at the same time:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).

User Interface Options	Path
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).
From the Connectors page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> (left panel).

- 3 Select the connectors whose destination runtime parameters you want to edit.
- 4 Click **Runtime Parameters** to open the wizard.
- 5 Follow these steps in the wizard to edit the runtime parameters:
 - a Select the destinations whose runtime parameters you want to modify.
 - b Select the configurations to be affected (default or alternate configurations).
 - c Select the group of parameters you want to modify (for example, batching, cache, network, processing).
 - d Modify the parameters.

Managing 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 pm time range and another configuration for the 5 pm to 8 am time range.

By default, a configuration labeled **Default** exists and is applied to a destination. Any subsequent configurations you define are labeled **Alternate#1**, **Alternate#2**, and so on. The default configuration is used if the time ranges specified for other alternate configurations do not span 24 hours. For example, if you specify an alternate configuration, **Alternate#1** that is effective from 7 am to 8 pm, the **Default** configuration will be used from 8 pm to 7 am (assuming that there are no other alternate configurations defined on this system).

If you need to apply the same alternate configuration for multiple destinations, you need to define an alternate configuration (with the same settings) for each of those destinations.

Defining a New Alternate Configuration


The process of defining a new alternate configuration includes first defining the configuration, and then editing it to specify the time range for which that configuration is effective.



To define an alternate configuration:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click () in the Destinations section.
- 4 Click **Add**.
- 5 Specify or update values for the listed parameters.
- 6 Scroll down to the end of the page and click **Save**.

If this is the first alternate configuration you defined, it is saved as Alternate#1. Subsequent configurations are saved as Alternate#2, Alternate#3, and so on.

To specify the time range for which the configuration you just defined is effective, edit the configuration you just defined using the following procedure [Editing an Alternate Configuration](#) below.

Editing an Alternate Configuration



In addition to editing an alternate configuration to change parameter values, you can edit it to specify the time range for which it is effective.

To edit an alternate configuration:

- 1 Click **Configuration > Manage Connectors**.

- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click () in the Destinations section.
- 4 Select the alternate configuration that you want to edit and click ().
- 5 Specify or update values for the listed parameters, including the time range in the From Hour/To Hour.
- 6 Scroll down to the end of the page and click **Save**.

Specifying a Time Range for an Alternate Configuration

See [“Editing an Alternate Configuration” on page 491](#).

Editing Alternate Configurations in Bulk

If you need to update the same parameters in multiple alternate configurations, follow the procedure described in [“Editing Destination Runtime Parameters” on page 489](#).

Sending a Command to a Destination


You can send a command to a connector destination.

To send a command to a destination on a connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).

User Interface Options	Path
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click  in the Destinations section.
- 4 Select the command you want to run and click **Next**.
- 5 Enter values for the parameters that the user interface displays and click **Finish**.

Removing a Connector



After removing a connector, you need to reboot the system; otherwise, the removed connector continues to forward events to its destination.


To remove a Connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel).

- 3 Select the connectors you want to delete. You can select multiple connectors.
- 4 Click **Delete** at the bottom of the page.
- 5 Reboot the system.



You can also delete a specific connector from its details page: Click **System** (left panel) > **Location** (left panel) > **Host** (left panel) > **Container** > **Connector** >  at the top of the page.


Sending a Command to a Connector


You can send a command to a connector.

To send a command to a connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click  in the Action column for the connector.

If you are on a specific Connector page,  is on top of the page.

- 4 From the **Command Type** drop-down list, select the command you want to send to the connector.
- 5 Click **Next**.

Running Logfu on a Connector


Run Logfu on a connector to parse ArcSight logs and generate an interactive visual representation of the information contained within the logs.

To run Logfu on a connector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).

User Interface Options	Path
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click () on top of the page. A separate window displays.

The system proceeds to retrieve and analyze system data logs. After this process is complete, a group of panels appears in the window.
- 4 From the **Group** box, choose which type of data you would like to view. The Group box lists all connectors within the chosen container, plus many other types of data such as memory usage, and transport rates and logs.

Choose one of the Group box **data points**. Depending on which data point you choose, a list of fields appears in the Field box below.
- 5 Choose a **field** to view. A graphic chart appears in the Chart box, providing rate and time information. The key at the bottom of the Chart box defines the data points mapped in the chart.
- 6 If you need to choose a different data point for analysis, click **Reset Data**.

Changing the Network Interface Address for Events

Connector Appliance has multiple network interfaces. By default, the connector determines which network interface address is used for events displayed in the ArcSight Console or Logger, but typically uses `eth0`.

To use a specific network interface address for events, add the parameter `connector.network.interface.name` to the Connector's `agent.properties` file. For example, to use the IP address for `eth1`, specify the following parameter:

```
connector.network.interface.name=eth1
```

Developing FlexConnectors

FlexConnectors are custom SmartConnectors that can read and parse information from third-party devices and map that information to ArcSight's event schema.

Connector Appliance provides a FlexConnector Development wizard that lets you quickly and easily develop a FlexConnector by creating a parser file, and enables you to test and package your new FlexConnector before deploying it. The wizard generates regular expressions and provides event field mapping suggestions automatically so you do not need to be an expert in regular expression authoring, parser syntax, or ArcSight event schema.

Use the FlexConnector Development wizard to develop FlexConnectors for simple log files. For complex log files, use the FlexConnector SDK (available from the HP Customer Support site (SSO)).



Note

Currently, the FlexConnector Development wizard supports Regex Files, Folder Follower, and Syslog (Daemon, File, Pipe) FlexConnectors only.

The FlexConnector Development wizard does not support the extra processors property or multiple sub messages. If you need these features, use the FlexConnector SDK to create your FlexConnector.




Caution

A FlexConnector that you develop with the FlexConnector Development wizard might perform more slowly than an ArcSight SmartConnector.

To develop a FlexConnector:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths to go to the **Containers** tab:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 Click  in the Action column of the container to which you want to add the FlexConnector. When the FlexConnector Development wizard opens, click **Next**.
- 4 Provide the vendor and product name of the device for which you are creating a FlexConnector, then click **Next**.



Note

The device vendor and product name are required.

- 5 Select the data source type, then click **Next**:
 - ◆ Select **Syslog** to create a Syslog FlexConnector to read events from Syslog messages.
 - ◆ Select **File** to create a FlexConnector to parse variable-format log files using regular expressions (ArcSight FlexConnector Regex File) or to parse variable-format log files in batch mode (ArcSight FlexConnector Folder Follower).
- 6 Upload a sample log file for the data source type you selected in the previous step, then click **Next**.
- 7 The wizard finds the first unparsed line in the log file, generates a regular expression to match and extract tokens from that line, and displays the suggested field mappings for each extracted token in the Mappings table.

Wizard

FlexConnector Development Wizard

Enter regular expression corresponding to text Lines Skipped: 0% Lines Parsed: 0%
Text 2005 Aug 24 13:57:54 EDT -04:00 %SPANTREE-6-PORTFWD: Port 3/16 state in VLAN 203 changed to forwarding

Regex Recalculate Reset

Mappings table

	Extracted Value	Type	Format	Event Field
1	2005 Aug 24 13:57:54	TimeStamp	yyyy MMM dd HH:mm:	deviceReceiptTime
2	3/16	String	String	deviceInboundInterface
3	203	Integer	String	deviceInboundInterface

Extra Mappings table

Event Field	Value
name	__stringConstant(SPAN)

Add Row

Cancel Skip Line Skip To End Previous Next



Note

The mappings are displayed in descending order of probability (based on ArcSight training data). You can change the mappings by selecting from the list.

The percentage of parsed lines in the file is shown in the top right of the panel. You can use this percentage to estimate where you are in the log file. The percentage of unparsed lines skipped in the file is also shown in the top right of the panel.

- ◆ To change the regular expression in the **Regex** box and recalculate the mappings, edit the expression and then click the **Recalculate** button. You can set the regular expression back to the suggested value by clicking the **Reset** button.
- ◆ Field mappings that do not correspond directly to the extracted tokens in the unparsed line of the log file are displayed in the Extra Mappings table. You can change the Event Field and provide a token operation. To add a new Event Field, click **Add Row**.

You can use extra mappings to:

- Remap an extracted token to a different Event Field in addition to the existing mapping. For example, you can add an Event Field with the value \$3 where \$3 is the third token in the list of suggested mappings.
- Map a modified token or combination of tokens to an Event Field. For example, you can add an Event Field with the value `__operation($1,$3)`.
- Map an Event Field to a constant string or integer. For example, you can add an Event Field with the value `__stringConstant(constant)`.



Note

The wizard always contains an extra mapping for the Event Field **name**, which maps all the words in the input log line. ArcSight strongly recommends that you do not simply delete the **name** Event Field but map it in either the Mappings or the Extra Mappings table.

For a list of the token operations used when tokens are mapped to ArcSight event fields, refer to the FlexConnector Developer's Guide (available from the HP Customer Support site (SSO)).

- 8 Click **Next** to save the mapping to the parser file and display the next unparsed line in the log file.



Click the **Skip Line** button to go to the next unparsed line in the log file without saving the mapping.

Click the **Skip to End** button to go to the end of the log file without processing any other lines and display the parser file for review.

Click the **Previous** button to go back to the previous line in the log file and make changes if necessary. If you configured any mappings for the previous line, the **Previous** button displays the configured mappings, not the default mappings.

After all unparsed lines in the log file have corresponding regular expressions and mappings, the wizard displays the parser file for review.

- 9 Review the parser file and make changes, if necessary, directly in the Review Parser File panel.



In Mozilla Firefox, if certain text in the Review Parser File panel is underlined in red, you can disable Spell Check; Right-click in the panel and click **Check Spelling** to remove the check mark.

- 10 Click **Next** to save and package the parser file.

- 11 Choose how you want to deploy the FlexConnector:

- ◆ Select **Deploy parser to existing connector in container** and click **Next** to use the parser file with an existing connector. Click **Done** to close the FlexConnector wizard and redisplay the Container tab.



The **Deploy parser to existing connector in container** option displays only if the container already contains a connector of the same type.

- ◆ Select **Add new connector to container** and click **Next** to add the parser as a new connector. Follow the steps to add the connector to the container.



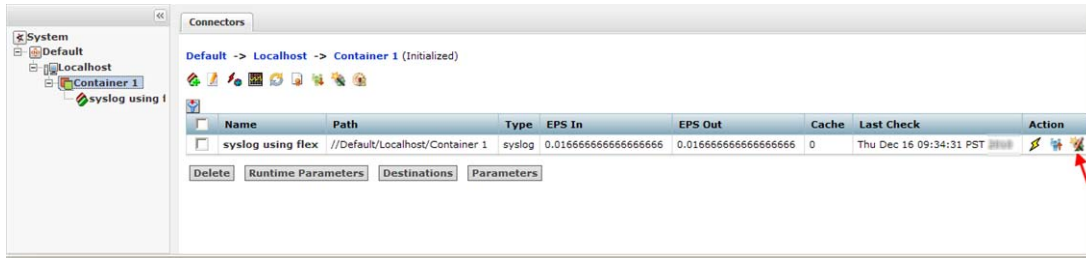
After deploying your FlexConnector, you can edit it any time from the **Connectors** tab. See [“Editing FlexConnectors” on page 498](#).


You can share FlexConnectors with other users. See [“Sharing Connectors \(ArcExchange\)” on page 499](#).

Editing FlexConnectors

After you have developed a FlexConnector with the FlexConnector wizard and have deployed it in a container, you can edit the FlexConnector to make changes to the parser file when needed.

The FlexConnector Edit wizard is available on the **Connectors** tab in the **Action** column.



Click  in the **Action** column for the FlexConnector to open the wizard. To edit the parser file, follow [Step 6](#) through [Step 11](#) in “[Developing FlexConnectors](#)” on [page 495](#).



Caution

Only edit a FlexConnector that is created with the FlexConnector wizard. Editing manually-created FlexConnectors might produce unpredictable results.



Tip

In addition to the FlexConnector Edit wizard, you can also use the Edit a File action in the Container Diagnostics wizard to edit your FlexConnector. Refer to “[Running Diagnostics on a Container](#)” on [page 473](#).

Sharing Connectors (ArcExchange)

You can share FlexConnectors and parser overrides with other users.

A FlexConnector is a custom connector that you define to gather security events from log files, databases, and other software and devices. You can share the following FlexConnector types:

- Syslog FlexConnectors (to read events from syslog messages)
- Log File FlexConnectors (to read fixed-format log files)
- Regular Expression Log File FlexConnectors (to read variable-format log files)
- Regular Expression Folder Follower FlexConnectors (to read variable-format log files recursively in a folder)
- Regular Expression Multiple Folder Follower FlexConnectors (to read events in real time or batch mode from multiple folders)
- XML FlexConnectors (to read events recursively from XML-based files in a folder)

A parser override is a file provided by ArcSight used to resolve an issue with the parser for a specific connector, or to support a newer version of a supported device where the log file format changed slightly or new event types were added. You can share parser overrides for all connector types that use a parser.

To share a FlexConnector or parser override, you need to package and upload it to ArcExchange on the ArcSight online community (Protect 724) or to your local machine. You can also download a FlexConnector or parser override that you need from ArcExchange or from your local machine and add it to a container.

Packaging and Uploading Connectors

Before uploading your FlexConnector or parser override to Protect 724 or to your local computer, you need to package it into a zip file, (called an AUP package) using the upload wizard.

A FlexConnector AUP package contains the connector properties file, categorization file, connector parameters, and a manifest file with all the metadata on the package required for successful deployment. Metadata includes information about the AUP package, such as the package type, connector type, connector description, and so on. You can create only one AUP package per connector per device type. You can package a FlexConnector in Basic or Advanced mode. In **Basic** mode:

- The wizard packages the FlexConnector properties file automatically. If the wizard finds more than one properties file, you are prompted to select the file you want to package.
- The wizard packages the categorization file automatically *only* if it can be determined based on the device vendor and product information found in the properties file.
- The wizard does not package connector parameters. You are prompted to configure the connector when it is downloaded and deployed.

In **Advanced** mode:

- The wizard packages the FlexConnector properties file automatically. If the wizard finds more than one properties file, you are prompted to select the file you want to package. (This is same as Basic mode.)
- The wizard packages the categorization file automatically if it can be determined based on the device vendor and product information found in the properties file. If the categorization file cannot be determined, you are prompted to select the categorization file you want to package from the list of files found in the container.
- The wizard displays connector parameters so you can configure the parameters you want to display and set the default values you want to provide during connector deployment (download). The parameters you do not configure for display are preconfigured with the current values and will not be displayed during connector deployment.

A parser override package contains the parser override properties file and the manifest file only.

Follow the steps below to package and upload a FlexConnector or parser override.



Note



- To upload to ArcExchange, you must have a valid username and password for Protect 724.
- Make sure that you have configured Connector Appliance network settings under Setup > System Admin > Network and that the appliance can communicate with the Protect 724 server.

To package and upload a FlexConnector or parser override:

- 1 Click **Configuration > Manage Connectors**.
- 2 Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the location in which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).

User Interface Options	Path
From the host on which the connector exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Connectors tab (right panel) > <i>Name of the Connector</i> (right panel).
From the Connector page	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > <i>Container</i> > <i>Name of the Connector</i> (left panel).

- 3 Click  at the top of the Connector page to open the upload wizard. (From the Connectors page, select the connector in the right panel and click  in the **Action** column.)
- 4 Click **Next** and follow the steps in the wizard to:
 - a Select the type of AUP package you want to create for the selected connector.
Connector Appliance scans the container and displays the relevant files that can be packaged.
 - b For a FlexConnector, select **Basic** to create a default package or select **Advanced** to customize the package to meet your needs. For a description of Basic and Advanced mode, refer to [“Packaging and Uploading Connectors” on page 499](#).
 - c If the connector contains several properties files, you are prompted to select the properties file you want to package. Certain connectors, for example, syslog connectors, can have more than one parser override folder, in this case, you are prompted to select the folder you want to package.
 - d If you selected Advanced mode for a FlexConnector in [Step b](#) and the categorization file cannot be determined, you are prompted to select the categorization file you want to package from a list of files found in the container.



Note

Categorization files are not packaged for parser overrides.

- e If you selected Advanced mode for a FlexConnector in [Step b](#), select the configuration parameters you want to display when the connector is deployed and then provide default values for these parameters. Parameters you do not select are pre-configured with the current values.

If any advanced connector parameters were previously modified from their defaults, the wizard displays these parameters so that you can select which ones you want to be configured automatically during deployment.



Note

Configuration parameters are not displayed for parser overrides.
If the connector has table parameters, they are not displayed during packaging. However, when the connector is downloaded to a container, you will be prompted to provide values for all the table parameters.

- f Provide a description of the AUP package and instructions on how configure the device used by the connector.

- g** Provide the vendor, product, and version of the device used by the connector.

If the wizard can determine the vendor, product, and version of the device, the information is displayed in the fields provided. You can change the information to meet your needs.

- h** Upload the created AUP package to ArcExchange or to your local machine.



To upload the AUP package to ArcExchange, you must have a valid username and password for Protect 724.

Downloading Connectors

You can download a FlexConnector or parser override that is available from ArcExchange on Protect 724 or from your local computer. You download a FlexConnector or parser override directly to a container.

You can download only one FlexConnector per container using the download wizard. However, there is no limit to the number of parser overrides you can download to a container.



- When downloading a parser override to a container, the download wizard overwrites any existing parser override with the same name in the container without prompting for confirmation. To avoid overwriting an existing parser override, send a **Get Status** command to the existing parser override to check the parser information before you download a new parser override. For information on sending a Get Status command, refer to [“Sending a Command to a Connector” on page 494](#).
- ArcSight recommends that you back up the container to the Backup Files repository before downloading a connector or parser override so you can revert to the previous configuration if the download produces unexpected results.

Follow the steps below to download a FlexConnector or parser override to a container.


To download to ArcExchange, you must have a valid username and password for Protect 724. Also, make sure that you have configured Connector Appliance network settings under Setup > System Admin > Network and that the appliance can communicate with the Protect 724 server.

To download a FlexConnector or parser override:

- Click **Configuration > Manage Connectors**.
- Go to the **Containers** page. Use one of these navigation paths:

User Interface Options	Path
From the System-level page	Click System (left panel) > Containers tab (right panel).
From the location in which the container exists	Click System (left panel) > <i>Location</i> (left panel) > Containers tab (right panel).

User Interface Options	Path
From the host on which the container exists	Click System (left panel) > <i>Location</i> (left panel) > <i>Host</i> (left panel) > Containers tab (right panel).

- 3 In the right panel, select the container into which you want to download the connector, and then click  in the **Action** column to open the download wizard.
- 4 Click **Next** and follow the steps in the wizard to:
 - a Select whether you want to download the connector from ArcExchange on Protect 724 or from your local computer.
 - b Select the AUP package you want to download.

On Protect 724, you can search for a parser override or FlexConnector AUP package using a keyword or a combination of keywords.



Note

You can only download a parser override package to a container that has a connector of the same type as the package.

You can download only one FlexConnector per container using the download wizard. If the container already contains a FlexConnector of the same type as the one you want to download, you can replace the existing FlexConnector with the one you are downloading, but you cannot create a new one.

- c For a FlexConnector, provide connector configuration parameters, if needed.

Preconfigured and advanced parameters are deployed automatically with the values that were packaged; you are not prompted to configure these parameters. The configurable parameters are displayed with suggested defaults, which you can modify if necessary. The table parameters are displayed with no configured values, you have to provide the values manually, as needed.

- d Add or select a destination for the connector.

If you are downloading the connector to a container that has an existing connector of the same type, you are *not* prompted for a destination.

The wizard copies the properties and categorization files to the appropriate locations and also installs the zip file for the AUP package in the `user/agent/deployedaups` folder on the Connector Appliance to keep track of the deployment history.

After a successful download, the container is restarted automatically.



Note

To use memory efficiently, parser overrides for the Windows Unified connector only load when the first event is received.

Configuration Suggestions for Connector Types

The following table provides configuration suggestions for different types of connectors.

Connector Type	Effects of Limited Usage
Syslog connectors	<p>Due to the nature of UDP (the transport protocol typically 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 might fill and the operating system drop UDP messages.</p> <p>Note: ArcSight recommends that you do not use the Limit CPU Usage option with these connectors because of the possibility of event loss.</p>
SNMP connectors	<p>Similar to Syslog connectors, when the event rate is limited on SNMP connectors, they can potentially lose events. SNMP is also typically UDP-based and has the same issues as Syslog.</p>
Database connectors	<p>Because 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 will be lost, unless the database tables are truncated. After the event burst is over, the connector might eventually catch up with the database if the event rate does not exceed the configured limit.</p>
File connectors	<p>Similar to database connectors, file-based connectors <i>follow</i> files and limiting their event rates 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. The connectors might catch up if the event rate does not exceed the configured rate.</p>
Asset Scanner connectors	<p>All connectors on Connector Appliance run as a service (not as an application). Therefore, asset scanner connectors running on Connector Appliance are <i>not</i> supported in Interactive mode.</p> <p>To run the asset scanner connector in Interactive mode, install the connector on a standalone system and manage it as a software-based connector.</p>
Proprietary API connectors	<p>The behavior of these connectors depends on the particular API, (for example, OPSEC behaves differently than PostOffice and RDEP). But in most cases, there will be no event loss unless the internal buffers and queues of the API implementation fill up. These connectors work much like database or file connectors.</p>

Deploying FlexConnectors

FlexConnectors are custom connectors that are user-defined. Connector Appliance ships with several prototype FlexConnectors, including:

- ArcSight FlexConnector File
- ArcSight FlexConnector ID-based Database
- ArcSight FlexConnector Multiple Database
- ArcSight FlexConnector Regular Expression File
- ArcSight FlexConnector Regular Expression Folder File

- ArcSight FlexConnector Simple Network Management Protocol (SNMP)
- ArcSight FlexConnector Time-based Database
- ArcSight FlexConnector XML File

You can create and manage FlexConnectors using repositories. You can share FlexConnectors with other Connector Appliance users. Refer to [“Sharing Connectors \(ArcExchange\)” on page 499](#).

For more information, consult the *FlexConnector Developer's Guide*, available from customer support.

Configuring the Check Point OPSEC NG Connector

The Check Point FW-1/VPN-1 OPSEC NG connector can operate in clear channel or sslca mode.



Note

- This procedure is supported only for ArcSight connector release 4.6.2 or later.
- A hostname is called an Application Object Name on Check Point. A password is a Communication Activation Key on Check Point.

To configure a connector to operate in sslca mode

On the Check Point SmartDashboard:


- 1 Create an OPSEC Application Object using the Check Point SmartDashboard. You need to provide these parameters when creating the application object.

Parameter	Description
Name	A meaningful name for the application object you are creating; for example, ArcSightLea-1. This name is used to pull the OPSEC certificate in the system.
Host	The hostname of the system managing the connector.
Client Entities	Select LEA.
Secure Internal Communication	If a DN string is not present, initialize the communication by providing an activation key. The activation key is used when the certificate is pulled. This is the SIC Name. Click Communication > Initialize .

After the object is created, note down the following information, which you will need to provide when continuing configuration.

- ◆ SIC Name—DN string that you obtain after initializing communication as described below.
- ◆ SIC Entity Name—Double-click the Check Point Gateway name in the SmartDashboard to view its general properties. The SIC Entity Name is the SIC string configured in the general properties window.
- ◆ Check Point IP address or hostname.

2 Pull the Check Point certificate.



To do so, run the Pull OPSEC Certificate command on the container to which you will be adding the connector. For detailed information about running a command on a container, see [“Running a Command on a Container” on page 470](#). You need to provide this information when running the command:

Parameter	Description
Server hostname or IP address	The name or IP address of the Check Point server.
Application object name	The OPSEC Application object name you specified in the previous step. This parameter is case sensitive.
Password	The activation key you entered when creating the OPSEC application object in the previous step.

If the certificate is pulled successfully, a message similar to this is displayed:

```
OPSEC SIC name (CN=ArcSightLea-1,0=cpfw1..5ad8cn) was retrieved
and stored in /opt/arcsight/<container name>/current/user/agent
/checkpoint/<name>. Certificate was created successfully and
written to "/opt/arcsight/<container name>/current/user/agent
/checkpoint/ArcSightLea-1.opsec.p12".
```

Note down the OPSEC SIC Name (CN=ArcSightLea-1,0=cpfw1..5ad8cn in the above example) and the file name (ArcSightLea-1.opsec.p12 in the above example).



If the certificate is not pulled successfully, check to ensure that the Application object name you specified is correct (including the case) and the container on which you are running the command is up and running.

3 Install Policy on the LEA client for the Check Point Gateway using the SmartDashboard.

On the Connector Appliance:

4 Add a Check Point connector by following instructions described in [“Adding a Connector” on page 475](#). You need to provide the following information.

Parameters	Values to input
Type	Check Point FW-1/VPN-1 OPSEC NG
Connection Type	SSLCA

Parameters	Values to input
Connector	Server IP: The IP address of the Check Point server.
Table	Server Port: The port on the server that listens for SSLCA connections. Use the default value 18184.
Parameters	OPSEC SIC Name: The name you noted in Step 1 . OPSEC SSLCA File: The name you noted after pulling the certificate in Step 2 . OPSEC Entity SIC Name: The name you noted in Step 1 .

- 5 An error similar to the following is displayed.

```
-1:[X] Unable to connect to the Lea Server[10.0.101.185] -1:1
connection test failed!
```

Click the **Ignore warnings** checkbox. Click **Next**.

- 6 Continue to configure the rest of the connector. Go to [Step 6](#) in “[Adding a Connector](#)” on page 475.

Adding the MS SQL Server JDBC Driver

When you install and configure database connectors that use Microsoft SQL Server as the database, a JDBC driver is required. This driver does not ship pre-installed on the system; you need to install it before configuring database connectors on the appliance.

To install a JDBC Driver:


- 1 Download the MS SQL Server JDBC Driver to a computer that can access Connector Appliance. You can download the driver from Microsoft at:
<http://msdn.microsoft.com/en-us/sqlserver/aa937724>
- 2 Run the setup program to install the driver.
- 3 Follow the instructions in “[Uploading Files to a Repository](#)” on page 520 to add the `sqljdbc.jar` file.



Tip

The name of the `jar` file may be different from that of some JDBC driver versions. Different versions of the JDBC driver are required for different SQL Server database versions; be sure to use the correct driver for your database.

The new driver file is added to the repository, as shown in the following example.



After you have installed the JDBC driver, you need to upload the driver file to the containers that will contain the SQL Server database Connectors. Follow the instructions in “[Uploading a File from the Repository](#)” on page 522.

After the driver file has been uploaded to a container, follow the instructions in [“Adding a Connector” on page 475](#) to add a connector that requires a JDBC driver.

Adding the MySQL JDBC Driver

When you install and configure database connectors that use MySQL as the database, a JDBC driver is required. This driver does not ship pre-installed on the system; you need to install it before configuring database connectors on the appliance.

To Install a JDBC Driver:

- 1 Download the MySQL JDBC Driver to a computer that can access Connector Appliance. You can download the driver from:

<http://dev.mysql.com/downloads/connector/j/5.0.html>

- 2 Extract the driver.

- 3 Follow the instructions in [“Uploading Files to a Repository” on page 520](#) to add the `mysql-connector-java-x.x.x-bin.jar` file.

The new driver file is added to the repository, as shown in the following example.

JDBC Drivers

Settings for JDBC Drivers

Retrieve Container Files

Upload To Repository

Add Connector JDBC Driver File

Name	Last Modified	Size	State
mysql-connector-java- 5.0.8.zip	Fri Mar 02 11:41:38 GMT-08:00 2012	8.39 MB	Retrieved

After you have installed the JDBC driver, you need to upload the driver file to the containers that will contain the MySQL database Connectors. Follow the instructions in [“Uploading a File from the Repository” on page 522](#).

After the driver file has been uploaded to a container, follow the instructions in [“Adding a Connector” on page 475](#) to add a connector that requires a JDBC driver.

Chapter 10

Managing Repositories

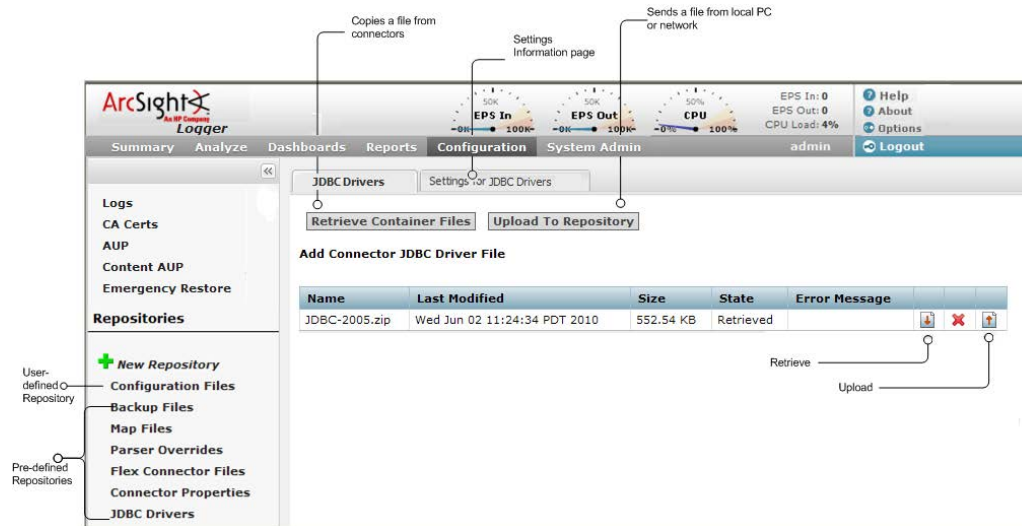
This chapter applies to Logger L3x00 appliance models only.

The following topics are discussed here.

- [“Overview” on page 509](#)
- [“Logs Repository” on page 511](#)
- [“CA Certs Repository” on page 511](#)
- [“UpgradeAUP Repository” on page 513](#)
- [“Content AUP Repository” on page 515](#)
- [“Remote Management AUP Repository” on page 516](#)
- [“Emergency Restore” on page 518](#)
- [“User-Defined Repositories” on page 518](#)
- [“Pre-Defined Repositories” on page 523](#)



Overview

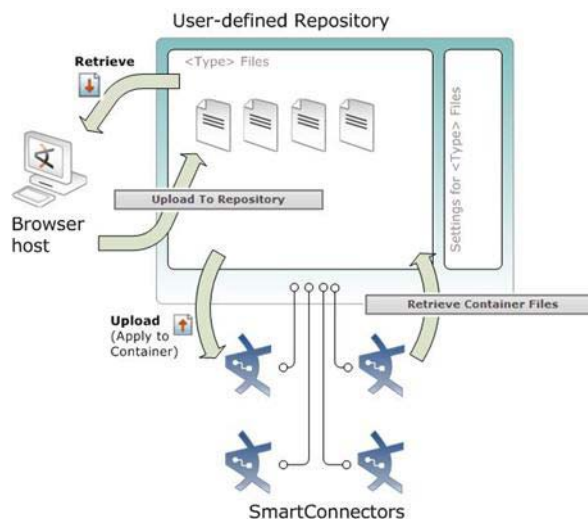
Certain management operations require a specific upgrade or content update (.enc) file, or a certificate. Other operations such as viewing the logs require you to load the logs to a Log repository. You can also maintain centralized repositories for files needed for connector configuration and management.



By default, a number of pre-defined repositories are provided. However, you can create more repositories to suit your needs. The repositories you create are referred to as user-defined repositories.

The following specific terms are used for repository functions.

- **Retrieve Container Files** copies a file from one or more connectors to the repository.
- **Upload to Repository** sends a file from your local computer (the computer running the browser) or a network host accessible from your local computer to the repository.
- **Retrieve**  downloads a file from the repository to your local computer network.
- **Upload**  copies a file from the repository to one or more connectors.



You can perform these operations using repositories:

- Manage logs in the Logs repository
- Manage CA certificates in the CA Certs repository
- Upgrade a connector using an upgrade file available in the Upgrade repository

- Apply a Content ArcSight Update Pack (AUP) on one or more connector
- Manage remote management configuration AUP files in the Remote Management AUP repository
- Restore a container when it is damaged and irrecoverable
- Maintain centralized repositories of files for connector configuration and management

Logs Repository

When you want to view connector logs, you need to first **Load** the logs of the container that contains the connector to the Logs repository, then **Retrieve** the logs to view them.



If a container contains more than one connector, logs for all connectors are retrieved.

For information on loading, retrieving, and deleting the logs, refer to the Connector Appliance Administrator's Guide.

Uploading a File to the Logs Repository

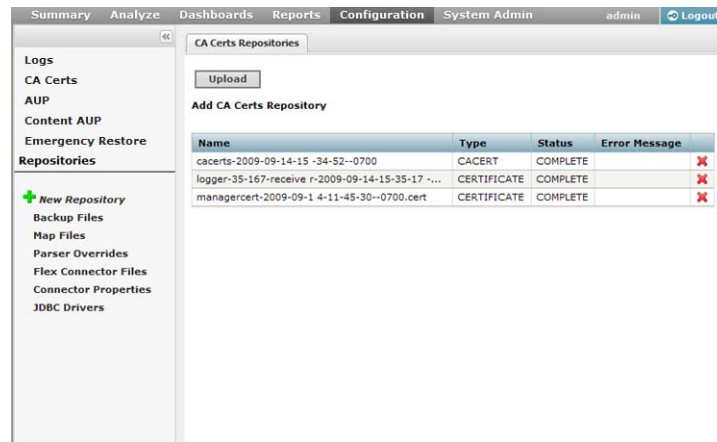
Uploading a file into the Log repository is useful for sharing annotated log or other files with other users. The file needs to be in .zip format.

To upload a file:

- 1 Click **Configuration > Repositories**.
- 2 Click **Logs** from the left panel.
- 3 Click **Upload** from the right panel.
- 4 Enter the local file path or click **Browse** to select the file.
- 5 Click **Submit** to add the specified file to the repository or **Cancel** to quit.

CA Certs Repository

Connectors require a Certificate Authority (CA) issued or self-signed SSL certificate to communicate securely with a destination. The CA Certs repository (shown below) enables you to store CA Certs files (that contain one or multiple certificates) and single CA certificates. When certificates are stored in the CA Certs repository, you can add the certificates to a container so that the connectors in the container can validate their configured destinations successfully.



To associate a CA certificate to a connector, you need to:

- Upload the CA certificate or CA Certs file to the CA Certs repository, as described below.
- Add a CA certificate from the CA Certs repository to the container that contains the connector, as described in the Connector Appliance Administrator's Guide.



You can add a single certificate to a container that is in FIPS or non-FIPS mode. You can only add a CA Certs file to a container that is in non-FIPS mode.

Uploading CA Certificates to the Repository

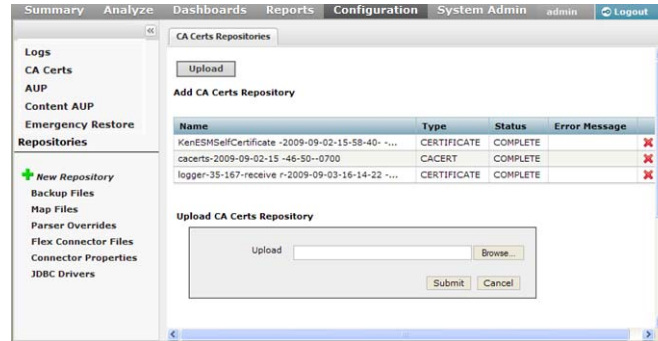
You can upload a CA Certs file or a single certificate to the CA Certs repository.



Before you upload a single CA certificate, change the name of the certificate on the local computer to a name that you can recognize easily. This helps you distinguish the certificate when it is displayed in the Certificate Management wizard.

To upload certificates to the repository:

- 1 Click **Configuration > Repositories**.
- 2 Click **CA Certs** in the left panel.
- 3 Click **Upload** in the right panel.
- 4 Enter the local path for the CA Certs file or the certificate, or click **Browse** to select it.
- 5 Click **Submit** to add the specified CA Certs file or the certificate to the repository, or **Cancel** to quit.



The CA Certs Repositories tab shows all the CA Certs files and single certificates that have been uploaded. The Type column shows CERTIFICATE for a single certificate and CACERT for a CA Certs file.

Removing CA Certificates from the Repository

You can delete a CA Certs file or a single certificate from the repository. When you delete a CA Certs file or a single certificate from the repository, it is deleted from the system.



Note

When you delete a CA Certs file or a single certificate from the CA Certs repository, containers are not affected; the connectors continue to use the certificates, which are located in a trust store after being added to a container. For information about adding a CA certificate to a container, refer to the Connector Appliance Administrator's Guide.

To remove a certificate from the repository:

- 1 Click **Configuration > Repositories**.
- 2 Click **CA Certs** in the left panel.
- 3 Identify the certificate or the CA Certs file you want to remove and click its associated Remove button (✖).

UpgradeAUP Repository

The Upgrade AUP repository enables you to maintain a number of connector AUP (upgrade) files. You can apply any of these AUP upgrade files to containers when you need to upgrade to a specific version. As a result, all connectors in a container are upgraded to the version you apply to the container.

This repository can also maintain upgrade files for upgrading remotely-managed Connector Appliances.

About the AUP Upgrade Process



The process discussed in this section only applies to upgrading connectors and to upgrading a remotely-managed Connector Appliance.

To upgrade a connector or to upgrade a remotely-managed Connector Appliance, you need to:

- Upload the appropriate .aup upgrade file to the Upgrade AUP repository, as described below.
- Apply the .aup upgrade file from the Upgrade AUP repository to the container or to a remote Connector Appliance. For more information, refer to the Connector Appliance Administrator's Guide.

Uploading an AUP Upgrade File to the Repository


To upload AUP upgrade files to the repository:

- 1 Download the upgrade AUP file for the connector or the remote Connector Appliance from the HP Customer Support site (SSO) at <http://support.openview.hp.com/> to the computer that you use to connect to the browser-based interface.
- 2 From the computer to which you downloaded the upgrade file, log in to the browser-based interface.
- 3 Click **Configuration > Repositories** from the top-level menu bar.
- 4 Click **UpgradeAUP** from the left panel.
- 5 Click **Upload** from the right panel.
- 6 Click **Browse** and select the file you downloaded earlier.
- 7 Click **Submit** to add the specified file to the repository or click **Cancel** to quit.
- 8 If you want to apply this upgrade file, follow these instructions:
 - ◆ For a container upgrade, see [“Upgrading a Container to a Specific Connector Version” on page 470](#).
 - ◆ For a remotely-managed Connector Appliance upgrade, see [“Upgrading a Host Remotely” on page 457](#).

Removing a Connector Upgrade from the Repository

You can remove a connector upgrade file from the repository when you no longer need it. When you remove a connector upgrade file from the repository, it is deleted from the system.

To remove a Connector upgrade from the repository:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 Click **UpgradeAUP** from the left panel.
- 3 Locate the upgrade file that you want to delete and click the associated  icon.

Content AUP Repository

ArcSight continuously develops new connector event categorization mappings, often called *content*. This content is packaged in ArcSight Update Packs (AUP) files. All existing content is included with major product releases, but it is possible to stay completely current by receiving up-to-date, regular content updates through ArcSight announcements and the HP Customer Support site (SSO). The AUP files are located under Content Subscription Downloads.

The ArcSight Content AUP feature enables you to apply an AUP file to applicable connector destinations that you are managing. Only the event categorization information can be applied to the connectors using this feature.

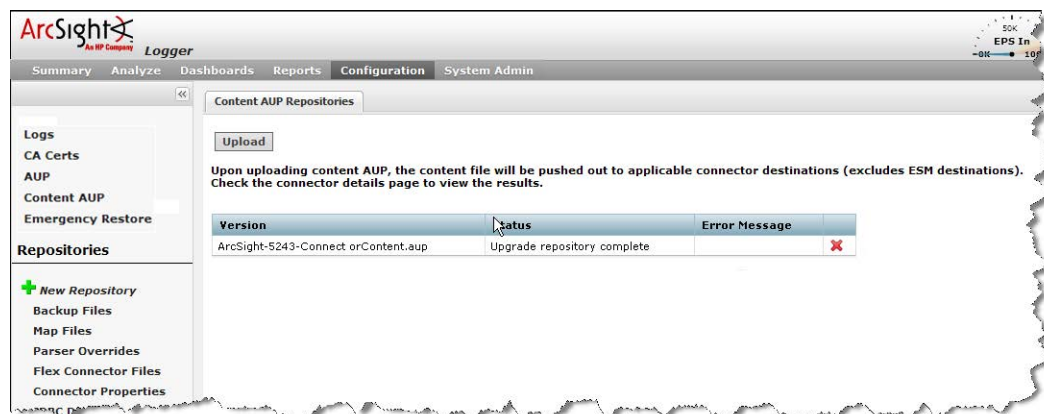
You can maintain a number of Content AUP files in the Content AUP repository. When an AUP file with a version number higher than the ones already in the repository is loaded, it is automatically pushed out to the connector destinations being managed. However, these connectors or connector destinations are skipped:

- Connectors that are unavailable at the time of the AUP file push
- Connectors whose current version does not fall in the range of versions that the Content AUP supports
- The ESM destination on a connector
- All destinations of a connector that have an ESM destination with the AUP Master flag set to Yes

Also, when a new connector is added, the highest number Content AUP is pushed automatically to its destinations.

Applying a New Content AUP

You can add a new content AUP file to the repository and push it automatically to all applicable connectors



To apply a new Content AUP:

- 1 Download the new Content AUP version from the HP Customer Support site (SSO) at <http://support.openview.hp.com/> to the computer that you use to connect to the browser-based interface.
- 2 From the computer to which you downloaded the AUP file, log in to the browser-based interface.

- 3 Click **Configuration > Repositories** from the top-level menu bar.
- 4 Click **Content AUP** from the left panel.
- 5 Click **Upload** from the right panel.
- 6 Click **Browse** and select the file you downloaded earlier.
- 7 Click **Submit** to add the specified file to the repository and push it automatically to all applicable connectors, or **Cancel** to quit.


You can verify the current Content AUP version on a connector by performing either of these steps:

- Run the `GetStatus` command on the connector destination and check that the value for `aup[acp].version` is the same as the AUP version you applied. For information about running a command on a connector destination, see [“Sending a Command to a Destination” on page 492](#).
- Hover your mouse over a connector name to see the AUP version applied to all destinations of that connector.

Applying an Older Content AUP

If you need to apply an older Content AUP from the Content AUP repository, delete all versions newer than the one you want to apply in the repository. The latest version (of the remaining AUP files) is pushed automatically to all applicable connectors.

To delete a Content AUP from the Content AUP repository:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 Click **Content AUP** from the left panel.
- 3 Locate the AUP file that you want to delete and click the associated  icon. Repeat for multiple files.

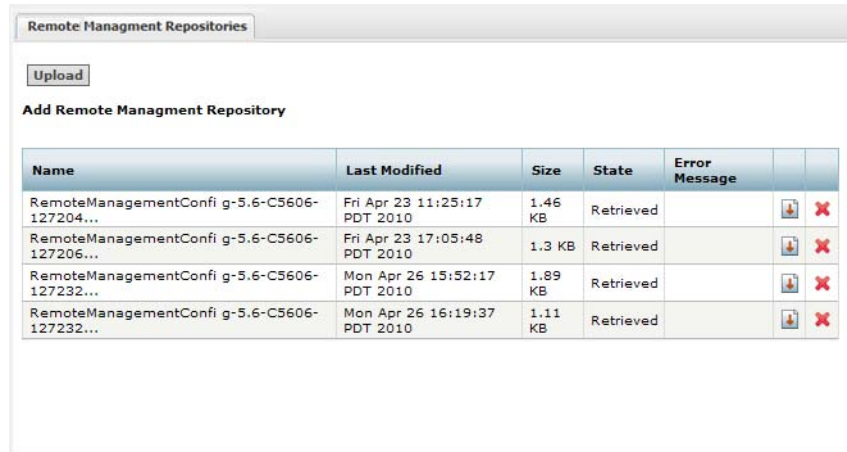
Remote Management AUP Repository

The Remote Management AUP repository stores AUP files that contain the remote management configuration of an appliance (a snapshot of all the remote software connectors and remote Connector Appliances that the appliance manages).

From the Remote Management AUP repository, you can:

- Download a Remote Management AUP file to your local computer (or network host accessible from the local computer) so that you can import the remote management configuration on another appliance.
- Upload Remote Management AUP files from your local computer (or network host accessible from the local computer) to the repository for storage.
- Delete Remote Management AUP files you no longer need.

The following example shows the Remote Management AUP repository.



Downloading Remote Management AUP Files

After you export the remote management configuration of a Connector Appliance, you can download the AUP file that contains the configuration to your local computer (or network host accessible from the local computer) so that it can be imported on another appliance.

For information on exporting and importing the remote management configuration of an appliance, refer to [“Exporting and Importing Remote Management Configuration” on page 449](#).

To download a Remote Management AUP file to your local computer:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 Click **Remote Management AUP** from the left panel.
- 3 Locate the AUP file in the table and click next to the file to download it to your local computer.

Uploading Remote Management AUP Files

You can upload remote management AUP files to the Remote Management AUP repository for storage.


To upload a Remote Management AUP file to the repository:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 Click **Remote Management AUP** from the left panel.
- 3 Click the **Upload** button at the top of the page.
- 4 Click **Browse** and select the file you want to upload from the local computer (or network host accessible from the local computer).
- 5 Click **Submit** to add the specified file to the repository.

Deleting Remote Management AUP Files

When a remote management AUP file is no longer up-to-date or needed, you can remove it from the repository.

To delete a Remote Management AUP file:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 Click **Remote Management AUP** from the left panel.
- 3 Locate the AUP file that you want to delete and click the associated  icon. Repeat for multiple files.

Emergency Restore

The Container Restore wizard guides you through the process of restoring a modified container. This feature is supported only for connectors and containers on the local host.


**Caution**

ArcSight recommends that you use this process only when a container is severely damaged and is no longer available. The Emergency Restore process deletes all information about that container and renders it empty. The connector is restored to the AUP version that you select.

**Note**

This feature is not available on the software version of Connector Appliance.

To restore a container:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 Click **Emergency Restore** from the left panel.
- 3 Follow the instructions in the Container Restore wizard.
- 4 Re-import the SSL certificate for the container. On the **Manage** tab, click the container name in the left panel. On the **Connectors** tab in the right panel, click the  icon to run the Certificate Download wizard and import the valid certificate.

User-Defined Repositories

A *user-defined repository* is a user-named collection of settings that control upload and download of particular files from connectors to the repository. Each repository uses a specified path, relative to `$ARCSIGHT_HOME/user/agent`, for files to be uploaded or for locations to download files. ArcSight connectors use a standard directory structure, so map files, for example, are always found in `$ARCSIGHT_HOME/user/agent`, (that is, the root directory, `$ARCSIGHT_HOME`, of the connector installation) in a folder called `map/`.

After they are created, user-defined repositories are listed on the left-side menu, under the **New Repository** heading, and appear with the user-specified display name.

User-defined repositories are expected to be grouped by file type and purpose, such as log files, certificate files, or map files. Each user-defined repository has a name, a display name, and an item display name, which are defined under the **Settings** tab that appears for user- or pre-defined repositories (for details about pre-defined repositories, see [“Pre-Defined Repositories” on page 523](#)).

Files viewed in the user-defined repository can be bulk processed with specified connectors and can be exchanged with the user's browser host.

Creating a User-Defined Repository

You can create a new repository at any time.



The repository requires correct directory paths. Your file will be applied to the wrong directory if the entered path contains errors, such as extra spaces or incorrect spellings. You can verify your directory paths by accessing the `Directory.txt` file, which lists the directory structure for every entered path. View the `Directory.txt` file by accessing your container logs and finding the `Directory.txt` file.

To create a new user-defined repository:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 Click **New Repository** under the Repositories section in the left panel.
- 3 For the new repository, enter the parameters listed in the following table.

Parameter	Description
Name	A unique name for the repository, typically based on the type of files it contains.
Display Name	The name that will be displayed on the left-side menu and for tabs: Process <i>names</i> , View <i>names</i> , Settings for <i>names</i> . Typically plural.
Item Display Name	The name used to describe a single item.
Recursive	Check to include sub-folders.
Sort Priority	-1 by default
Restart Connector Process	Check to restart the connector process after file operations.
Filename Prefix	An identifying word that is included in the names of retrieved files. For example, map files are identified by <code>map</code> in the file name: <code>localhost_Container_-1.Map-2009-04-06_12-22-25-607.zip</code>
Relative path (Download)	The path for download, relative to <code>\$ARCSIGHT_HOME</code> , for example, <code>user/agent/map</code> OR <code>user/agent/flexagent</code> . Leave this field blank to specify files in <code>\$ARCSIGHT_HOME</code> . Note: The relative path is used for download only.
Include Regular Expression	A description of filenames to include. Use <code>.*</code> to specify all files. The following example selects properties files that consist of <code>map</code> , followed by one or more digits, followed by <code>.properties</code> : <code>map\[0-9]+\\.properties\$</code>
Exclude Regular Expression	A description of filenames to exclude. The following example excludes all files with a certain prefix or in the <code>agentdata</code> folder. <code>(agentdata/ cwsapi_fileset_).*\$</code>
Delete Before Upload	Check to delete earlier copies before upload. CAUTION: If you check Delete Before Upload and do not specify a Relative path (Upload), all files and folders in <code>current/user/agent</code> will be deleted.

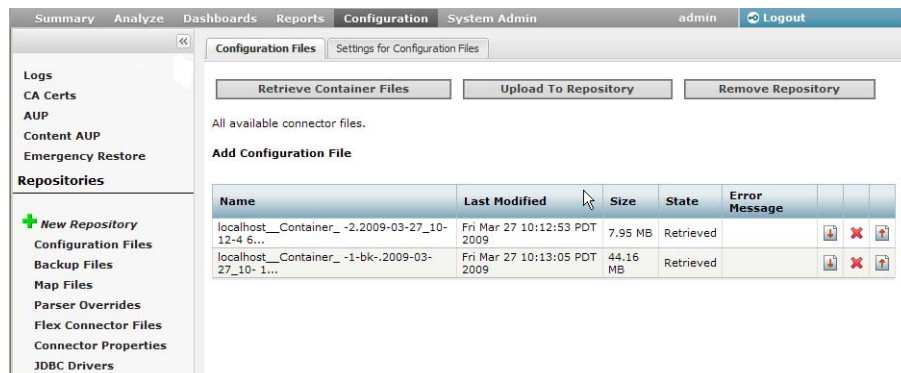
Parameter	Description
Delete Groups	Whether to delete folders recursively in \$ARCSIGHT_HOME/user/agent/map directory.
Relative path (Upload)	The path for upload, relative to \$ARCSIGHT_HOME/current/user/agent/flexagent/<connectorname>
Delete Relative Path	Whether the directory specified in Relative Path (Upload) and its contents should be removed when a file is uploaded from the repository.
Delete Include Regular Expression	Typically the same as the Include Regular Expression.
Delete Exclude Regular Expression	Typically the same as the Exclude Regular Expression.

- 4 Click **Save** at the bottom of the page.

The new repository displays under the **New Repository** heading in the left-side window panel.

Retrieving Container Files

The Retrieve Container Files button copies a file from one or more connectors to a repository. The specific files that are retrieved depend on the settings of a repository.



To retrieve a container file:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 In the left panel, click the name of the repository to which you want to copy connector files.
- 3 Click **Retrieve Container Files** in the right panel.
- 4 Follow the instructions in the Retrieve Container Files wizard.

Uploading Files to a Repository

The upload process copies files from your local computer to a repository.

To upload files to a repository:

- 1 Click **Configuration > Repositories** from the top-level menu bar.

- 2 In the lower left panel (under **Repositories**), click the name of the repository to which you want to upload files.
- 3 Click **Upload To Repository** from the right panel.
- 4 Follow the instructions in the Repository File Creation wizard.

Although you can select Repository zip file in the **Select the type of file that you want to upload** page of the Repository File Creation wizard, ArcSight recommends that you select **Individual files** to create a zip file with appropriate path information.

Be sure **not** to change the default sub-folder name `lib` in the **Enter the sub folder where the files will be uploaded** page of the Repository File Creation wizard.

Deleting a Repository

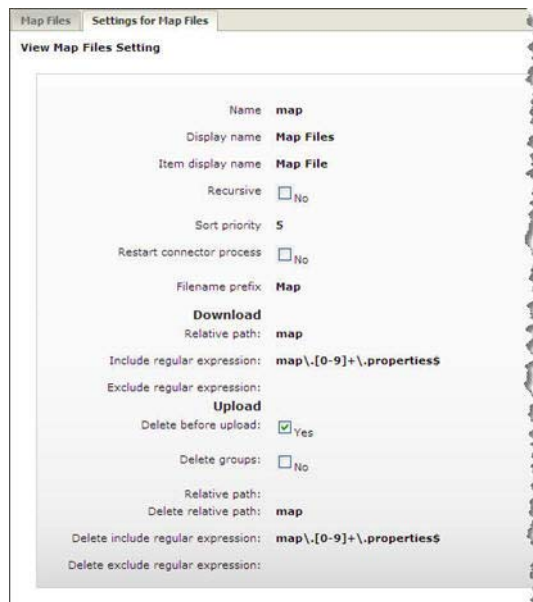
You can delete user-defined repositories only.

To delete a repository:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 From the left panel, click the name of the repository you want to delete.
- 3 Click **Remove Repository** from the right panel.

Updating Repository Settings

The Settings tab displays the settings associated with the current repository. An example is shown below. Most settings for pre-defined repositories are read-only; however, you can update settings for user-defined repositories.



To update settings of a repository:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 In the left panel, click the name of the repository whose settings you want to update.
- 3 Click the **Settings for *Repository_Name*** tab from the right panel.

- 4 Update the settings.
- 5 Click **Save** at the bottom of the page.

Managing Files in a Repository

You can retrieve files in a repository (download files to your local computer network), upload files to a repository, or remove files from a repository.



Caution

Connectors require correct properties and proper files. Applying incorrect files, including empty files or files with binary content, can prevent a connector from functioning correctly.




Tip

It is possible to upload files with incorrect content, such as an empty `.map` file. The system does not check or warn against such files. To ensure a successful result, only upload known, correct files.


Retrieving a File from the Repository

To retrieve a file from the repository:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 From the left panel, click the name of the repository in which the file exists.
- 3 Click  from the right panel for the file that you want to retrieve.
- 4 Follow the file download instructions to copy the file to your local computer.


Uploading a File from the Repository

To upload a file from the repository:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 In the left panel, click the name of the repository in which the file exists.
- 3 In the right panel, click  next for the file that you want to upload.
- 4 Follow the Upload Container Files wizard instructions to upload the file to the containers of your choice.
- 5 Verify that the file was uploaded correctly:
 - ◆ If you have SSH access to the connectors, connect to them and check the file structure.
 - ◆ Obtain the connector logs and check the contents of the `Directory.txt` file for each connector.

Removing a File from the Repository

To remove a file from the repository:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 In the left panel, click the name of the repository in which the file exists.
- 3 In the right panel, click  for the file that you want to delete.

Pre-Defined Repositories

You can define repositories for any connector-related files. As a convenience, the following repositories are pre-defined.

- **Backup Files:** connector cloning (see [“Cloning Container Configuration” on page 526](#)).
- **Map Files:** enrich event data
- **Parser Overrides:** customize the parser (see [“Adding Parser Overrides” on page 527](#))
- **Flex Connector Files:** user-designed connector deployment
- **Connector Properties:** `agent.properties`; subset of cloning
- **JDBC Drivers:** database connectors

To view the settings for a pre-defined repository, click the name of the repository and then click the **Settings** tab in the right panel.



Note

The settings for pre-defined repositories are read-only; to modify the settings, click **New Repository** in the left panel to create a user-defined repository and provide the settings you want to use.

The following tables lists the settings for each pre-defined repository.

Settings for Backup Files

Name	Default Setting
Name	backup
Display Name	Backup Files
Item Display Name	Backup File
Recursive	Checked (Yes)
Sort Priority	0
Restart Connector Process	Checked (Yes)
Filename Prefix	ConnectorBackup
Download Relative Path	
Download Include regular expression	
Download Exclude regular expression	(agentdata/ cwsapi_fileset_).*\$
Delete before upload	Checked (Yes)
Delete groups	Checked (Yes)
Upload Relative Path	
Delete Relative Path	
Delete Include regular expression	
Delete Exclude regular expression	(agentdata/ cwsapi_fileset_).*\$

Settings for Map Files

Name	Default Setting
Name	map
Display Name	Map Files
Item Display Name	Map File
Recursive	Un-checked (No)
Sort Priority	5
Restart Connector Process	Un-checked (No)
Filename Prefix	Map
Download Relative Path	map
Download Include regular expression	map\[0-9]+\properties\$
Download Exclude regular expression	
Delete before upload	Checked (Yes)
Delete groups	Un-checked (No)
Upload Relative Path	
Delete Relative Path	map
Delete Include regular expression	map\[0-9]+\properties\$
Delete Exclude regular expression	

Settings for Parser Overrides

Name	Default Setting
Name	parseroverrides
Display Name	Parser Overrides
Item Display Name	Parser Override
Recursive	Checked (Yes)
Sort Priority	10
Restart Connector Process	Checked (Yes)
Filename Prefix	Parsers
Download Relative Path	fcv
Download Include regular expression	.*
Download Exclude regular expression	
Delete before upload	Checked (Yes)
Delete groups	Checked (Yes)
Upload Relative Path	

Name	Default Setting
Delete Relative Path	fcp
Delete Include regular expression	.*
Delete Exclude regular expression	

Settings for FlexConnector Files

Name	Default Setting
Name	flexconnectors
Display Name	Flex Connector Files
Item Display Name	Flex Connector File
Recursive	Checked (Yes)
Sort Priority	15
Restart Connector Process	Checked (Yes)
Filename Prefix	FlexConnector
Download Relative Path	flexagent
Download Include regular expression	.*
Download Exclude regular expression	
Delete before upload	Checked (Yes)
Delete groups	Checked (Yes)
Upload Relative Path	
Delete Relative Path	flexagent
Delete Include regular expression	.*
Delete Exclude regular expression	

Settings for Connector Properties

Name	Default Setting
Name	connectorproperties
Display Name	Connector Properties
Item Display Name	Connector Property File
Recursive	Un-checked (No)
Sort Priority	20
Restart Connector Process	Checked (Yes)
Filename Prefix	ConnectorProperties
Download Relative Path	

Name	Default Setting
Download Include regular expression	agent\..*
Download Exclude regular expression	
Delete before upload	Un-checked (No)
Delete groups	Un-checked (No)
Upload Relative Path	
Delete Relative Path	
Delete Include regular expression	agent\..*
Delete Exclude regular expression	

Settings for JDBC Drivers

Name	Default Setting
Name	jdbcd drivers
Display Name	JDBC Drivers
Item Display Name	Connector JDBC Driver File
Recursive	Un-checked (No)
Sort Priority	25
Restart Connector Process	Checked (Yes)
Filename Prefix	
Download Relative Path	lib
Download Include regular expression	
Download Exclude regular expression	
Delete before upload	Un-checked (No)
Delete groups	Un-checked (No)
Upload Relative Path	
Delete Relative Path	lib
Delete Include regular expression	
Delete Exclude regular expression	

Cloning Container Configuration

Using the **Backup Files** repository, you can quickly copy a container to other containers. As a result, all connectors in the source container are copied to the destination container. This process is called *cloning* a container configuration. You can clone a container to

several containers at once. The contents of the source container replace the existing contents of the destination container.



Caution

Do not clone older, software-based connectors (such as build 4.0.8.4964) to containers with newer connector builds (such as 4.0.8.4976 or later).

Cloning a connector using the Backup repository only works if the connector version numbers are the same.

To clone a container:

- 1 Click **Configuration > Manage Connectors** to list the containers and determine the source and destination for cloning.
- 2 Click **Configuration > Repositories** from the top-level menu bar.
- 3 Click **Backup Files** under the **Repositories** section in the right panel.
- 4 If the backup file that you need to use for cloning exists in the repository, go to the next step. Otherwise, follow the instructions in [“Retrieving a File from the Repository” on page 522](#) to retrieve the container’s backup file to the Backup repository.

The retrieved file is named in `<connector name> ConnectorBackup <date>` format.

- 5 Follow the instructions in [“Uploading a File from the Repository” on page 522](#) to upload the backup file to one or more containers.

The destination containers are unavailable while the backup file is applied and the connectors are restarted.



Note

The backup file does not include the container certificates. You have to re-apply the certificates to the container after you upload the backup file.

After applying the certificates, check the status of the destination container to make sure it is available.

Adding Parser Overrides

A parser override is a file provided by ArcSight used to resolve an issue with the parser for a specific connector, or to support a newer version of a supported device where the log file format changed slightly or new event types were added.

To use parser overrides, you need to:

- Upload a parser override file to the pre-defined **Parser Overrides** repository.
- Download the parser override file to the container that contains the connector that will use the parser override.

Follow the steps below.

To upload a parser override file:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 Click **Parser Overrides** under the **Repositories** section in the right panel.
- 3 On the **Parser Overrides** tab, click the **Upload To Repository** button.
- 4 Follow the wizard to upload the file. When prompted by the wizard, make sure you:
 - ◆ Select the **Individual Files** option from the **Select the type of file that you want to upload** field.

- ◆ Add a slash (/) after fcp before adding the folder name in the **Enter the sub folder where the files will be uploaded** field. For example, `fcp/multisqlserver_audit_db`.

When upload is complete, the parser override file is listed in the table on the Parser Overrides tab.

To download the parser override file to a container:

- 1 Click **Configuration > Repositories** from the top-level menu bar.
- 2 Click **Parser Overrides** under the **Repositories** section in the right panel.
- 3 In the table on the **Parser Overrides** tab, locate the parser override file you want to download and click the up arrow next to the file.
- 4 Follow the wizard to select the container to which you want to add the parser overrides.

When the wizard completes, the parser overrides will be deployed in the selected container.



You can download a parser override file from ArcExchange. For more information, refer to [“Sharing Connectors \(ArcExchange\)” on page 499](#).

To verify that the parser override has been applied successfully, issue a Get Status command to the connector. See [“Sending a Command to a Destination” on page 492](#). In the report that appears, check for the line starting with the text `ContentInputStreamOverrides`.

Appendix A

Search Operators

This appendix describes the operators you can use in search queries you specify in the Search box (**Analyze > Search**) and gives examples of their use.

This appendix provides information on the following search operators.

["cef \(Deprecated\)" on page 529](#)

["chart" on page 530](#)

["dedup" on page 536](#)

["eval" on page 536](#)

["extract" on page 537](#)

["fields" on page 539](#)

["head" on page 539](#)

["keys" on page 540](#)

["parse" on page 541](#)

["rare" on page 542](#)

["regex" on page 542](#)

["rename" on page 543](#)

["replace" on page 544](#)

["rex" on page 545](#)

["sort" on page 547](#)

["tail" on page 548](#)

["top" on page 548](#)

["transaction" on page 549](#)

["where" on page 551](#)

cef (Deprecated)

Prior to Logger 5.2, you needed to use the cef operator to extract CEF fields from CEF events that matched the indexed search filter (the query portion before the first pipeline in the query expression) before you could use other search operators to act upon those fields. However, starting with Logger 5.2, you do not need to explicitly extract the CEF fields and then apply other search operators to those fields. You can specify the event fields directly in queries.

Extracts values for specified fields from matching CEF events. If an event is non-CEF, the field value is set to NULL.

Usage

```
...| cef <field1> <field2> <field3> ...
```

Notes

If multiple fields are specified, separate each field name with a white space or a comma.

To identify the name of a CEF field, use the Search Builder tool (click Advanced Search under the Search text box), which lists the names of all fields alphabetically.

The extracted fields are displayed as additional columns in the All Fields view (of the System Fieldsets). To view only the extracted columns, select **User Defined Fieldsets** from the System Fieldsets list.

Example 1

```
...| cef categorySignificance agentType
```

Example 2

```
...| cef deviceEventCategory name
```

chart

Displays search results in a chart form of the specified fields.

Usage

```
...| chart <field>
```

```
...| chart count by <field1> <field2> <field3> ...  
[span [<time_field>]=<time_bucket>]
```

```
...| chart {{sum | avg | min | max | stdev} (<field>)}+ by <field1>,  
<field2>, <field3> ...[span [<time_field>]= <time_bucket>]
```

```
...| chart {<function> (<field>)} as <new_column_name> by <field>  
[span [<time_field>]=<time_bucket>]
```

where

<field>, <field1>, <field2> are the names of the field that you want to chart. The fields can be either event fields available in the Logger schema or a user-defined fields created using the rex or eval operator prior in the query.

<time> is the bucket size for grouping events. Use d for day, h for hour, m for minute, s for seconds. For example, 2h, 5d, 1m. (See Notes for details.)

<function> is one of these: count, sum, avg (or mean), min, max, stdev

<new_column_name> is the name you want to assign to the column in which the function's results are displayed. For example, Total.


Deprecated Usage

The following deprecated usage contains “_count”. The recommended usage, as shown above, is “count”.

```
...| chart _count by <field1> <field2> <field3> ...
```

Notes

By default, a column chart is displayed. Other chart types you can select from: bar chart, line chart, pie chart, area chart, stacked column, or stacked bar.

To change the chart settings (including its type), click  to the upper right corner of the Result Chart frame of the screen. You can change these settings:

- **Title:** Enter a meaningful title for the chart.
- **Type:** Column, Bar, Pie, Area, Line, Stacked column, Stacked Bar. The last two types create stacked charts in which multiple values are plotted in a stack form. These charts are an alternate way of representing multi-series charts, which are described below.
- **Display Limit:** Number of unique values to plot. Default: 10

If the configured Display Limit is less than the number of unique values for a query, the top values equal to the specified Display Limit are plotted. That is, if the Display Limit is 5 and 7 unique values are found, the top 5 values will be plotted.

All chart commands except “count by” accept only *one field* in the input. The specified field must contain numeric values.

If multiple fields are specified, separate the field names with a white space or a comma.

The chart <field> command does not aggregate field values. It simply lists and charts each occurrence of the values of the specified field. For example, `chart sourcePort`. However, when you use an **aggregation function** such as `count by`, `sum`, `avg` (or `mean`), and so on, an aggregation of the specified fields is performed and charted, as illustrated in “[Example 1](#)” on page 534.

You can click on a charted value to quickly filter down to events with specific field values. For more information, see “[Chart Drill Down](#)” on page 114.

Aggregation Functions



Aggregation functions can only be used on numeric fields. The specified fields must contain numeric values. If a field you specify is of the wrong data type, you will receive an error message like the following: “The search cannot be run, there is an error in your query: Invalid field type for field [field name].”

If an aggregation function such as `count`, `sum`, or `avg` is specified, a chart of the aggregated results is displayed along with the tabular results of the aggregation operation in a Results Table. For example, for the aggregation function `sum(deviceCustomNumber1)`, the `sum_deviceCustomNumber1` column in the Results Table displays the sum of unique values of the `deviceCustomNumber1` field.

If this field had two unique values 1 and 20, occurring 2 times each, the `sum_deviceCustomNumber1` column displays sum of those two values, as shown in the following figure:

Result Table	
Page 1 of 1	
deviceCustomLumber1	sum_deviceCustomLumber1
1	2
20	40

The mathematical operators `avg` and `mean` are identical.

You can include multiple functions in the same `chart` command. When doing so, separate each function with a comma, as shown in this example:

```
...| chart count, sum(deviceCustomNumber3) by deviceEventClassId
```

When you include multiple functions, one column per function is displayed in the search Results Table. The Results Chart, however, plots the chart for the field specified in the “by” clause.

You can use the “as new_column_name” clause to name any column resulting from the aggregation functions, as shown in this example:

```
...| chart sum(deviceCustomNumber3) as TotalStorage,  
avg(deviceCustomNumber3) as AverageStorage by deviceCustomNumber3
```

Once defined, the newly defined column can be used in the pipeline as any other field. For example,

```
...| chart sum(deviceCustomNumber3) as TotalStorage,  
avg(deviceCustomNumber3) as AverageStorage by deviceCustomNumber3 |  
eval UpdatedStorage = TotalStorage + 100
```

When you export the search results of a chart operator, the newly defined column name (using the `chart` function as `new_column_name` command) is preserved.

Multi-Series Charts

A multi-series chart can plot the values of multiple aggregation functions in a single chart.

If you include multiple aggregation functions in a `chart` command, Logger generates a multi-series chart that plots the values of the specified aggregation functions along the Y-axis, as illustrated in [“Example 2” on page 534](#). Multi-series charts can be any of the chart types except Pie charts. For example, you can choose to plot a multi-series chart as a stacked chart—Stacked column or Stacked Bar—in which multiple values are plotted in a stack form, as illustrated in [“Example 3” on page 535](#).


The span Function

In addition to grouping events by the Logger schema fields (or the ones defined by the `rex` or `eval` operators), the `span` function provides an additional way to group events by a time field (such as `EventTime` or `deviceReceiptTime`) and a time bucket. In the following example, `deviceReceiptTime` is the time field and `5m` (5 minutes) is the time bucket:

```
...| chart count by deviceEventCategory span (deviceReceiptTime) =  
5m
```


If a time field is not specified for the `span` function, `EventTime` is used as the default. For example, the following query uses `EventTime` by default:

```
...| chart count by deviceEventCategory span = 5m
```

By default, the `chart` command displays the first 10 unique values. If the `span` function creates more than 10 unique groups, not all of them will be displayed. If you want to view all of the unique groups, increase the Display Limit value under Chart Settings. (Click  to the upper right corner of the Result Chart frame of the screen.)

Grouping with `span` is useful in situations when you want to find out the number of occurrences in a specific time span.

If you want to find out the total number of incoming bytes every 5 minutes on a device, you can specify a span of 5m, as shown in this example:

```
...| chart sum(deviceCustomNumber1) span=5m
```

The above example assumes that `deviceCustomNumber1` field provides the incoming bytes information for these events.

The `span` field can be used for grouping in conjunction with or without the event fields that exist in Logger schema or user-defined fields using the `rex` or `eval` operators. When a `span` field is specified in conjunction with an event field, the unique sets of all those fields is used for grouping. The following example uses `deviceCustomNumber3` and `deviceAddress` in conjunction with `span` to find out the number of events (using `deviceCustomNumber3`) from a specific source (using `deviceAddress`) in one hour:

```
...| chart sum (deviceCustomNumber3) by deviceAddress span=1h
```

When `span` is included in a query, search results are grouped by the specified time bucket. For example, if `span=5m`, the search results will contain one row for each 5-minute span. If there are no events within a specific 5-minute span, that row will be empty.

Additionally, the `span` function assumes a 24-hour day, all year long. If `span=1d` or `24h`, on the day of daylight savings time change, the event time indicated by the `span_eventTime` field in the search results will be different from the previous day by one hour. On the day when there are 23 hours in a day (in March), the span bucket will still include events from the last 24 hours. Similarly, on the day when there are 25 hours in the day (in November), the span bucket will include events from the last 24 hours. The following example illustrates the `span_eventTime` field when the span time bucket is 1d and the daylight savings times occurs on

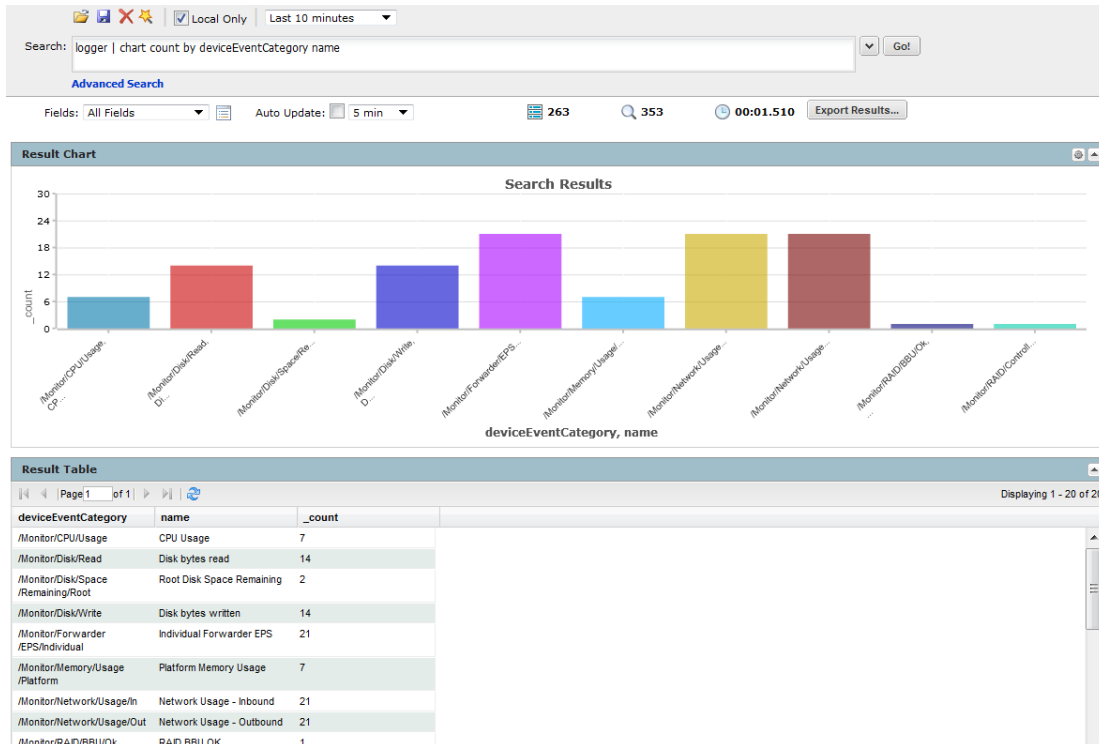
March 14th, 2011 and November 7, 2011:

span_eventTime	avg_logins
3/11/2011 12am	8
3/12/2011 12am	10
3/13/2011 12am	4
3/14/2011 1am	6
3/15/2011 1am	7
...	
11/5/2011 1am	4
11/6/2011 1am	2
11/7/2011 12am	5
11/8/2011 12am	7
...	

Example 1

Use the default chart setting (Column Chart) to specify multiple fields. In this example, a count of unique groups of deviceEventCategory and name fields is displayed and plotted.

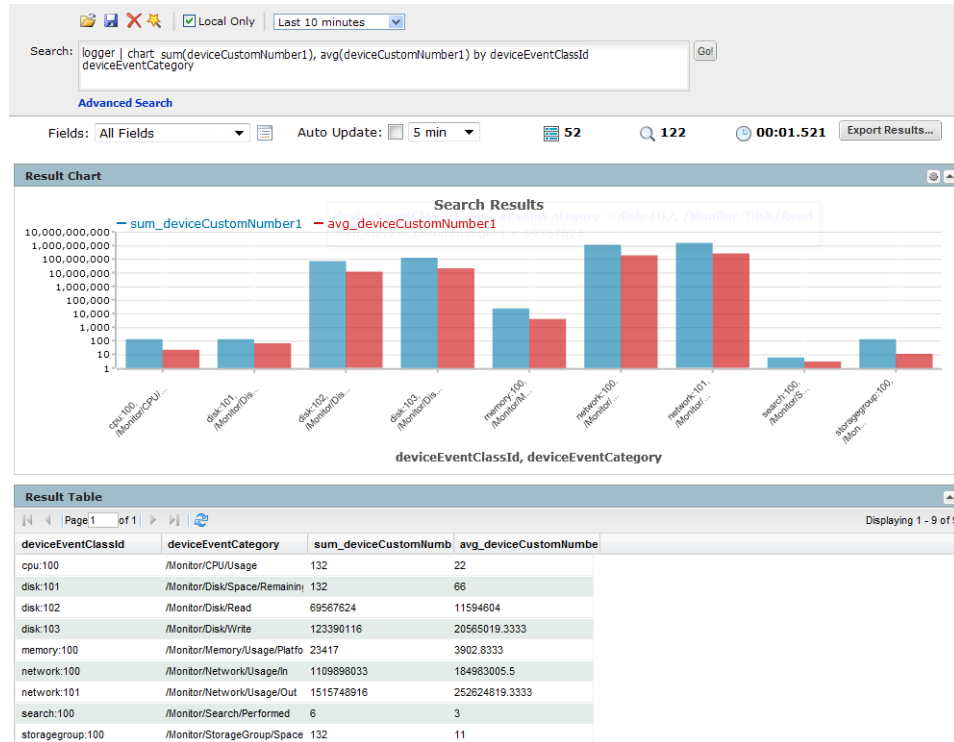
```
... | chart count by deviceEventCategory name
```

**Example 2**

Include average and sum in a chart command, to generate a multi-series chart that plots the values of these functions along the Y-axis in a single chart.

In the following query, unique groups of deviceEventClassId and deviceEventCategory are plotted along the X-axis, and the sum of deviceCustomNumber1 and average of deviceCustomNumber2 is plotted along the Y-axis.

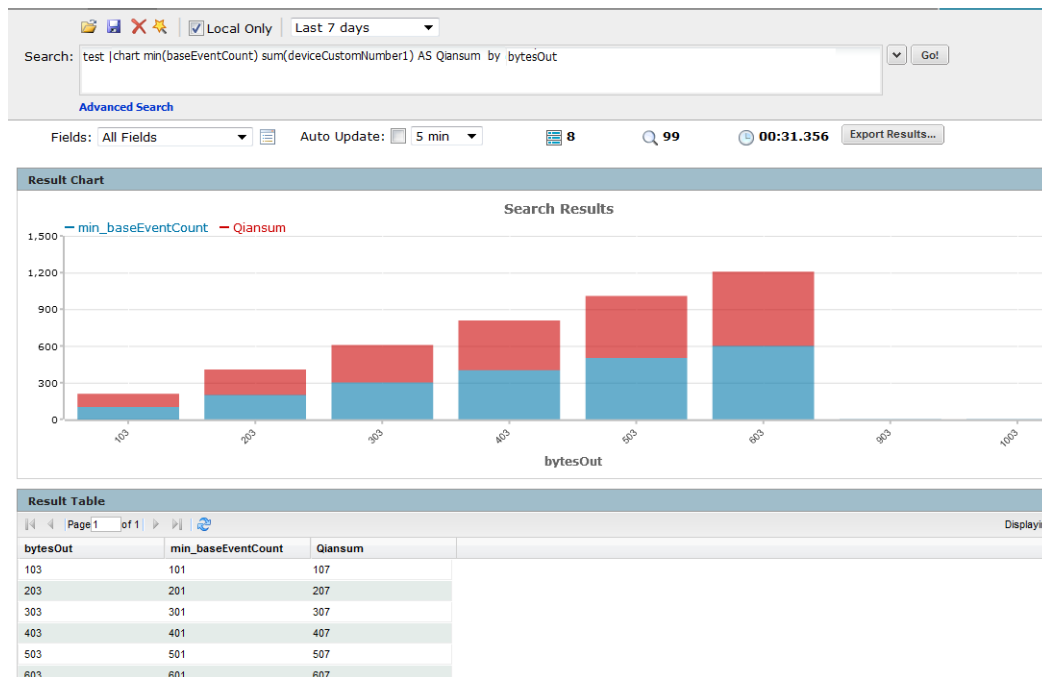
```
... | chart sum(deviceCustomNumber1), avg(deviceCustomNumber1) by deviceEventClassId deviceEventCategory
```



Example 3

Plot a multi-series chart as a stacked chart—Stacked column or Stacked Bar—in which multiple values are plotted in a stack form, as shown in the following figure.

```
... | chart min(baseEventCount) sum(deviceCustomerNumber1) AS Qiansum
by bytesOut
```



dedup

Removes duplicate events from search results. That is, events that contain the same value in the specified field. The first matching event is kept, and the subsequent events with the same value in the specified field are removed.

Usage

```
... | dedup [N] <field1>,<field2>, ... [keepevents=(true|false)]  
[keepempty=(true|false)]
```

`N` is an optional number that specifies the number of duplicate events to keep. For example, “dedup 5 deviceEventClassId” will keep the first five events containing the same deviceEventClassId values for each deviceEventClassId, and remove the events that match after the first five have been kept. Default: 1.

`field1, field2` is a field or a comma-separated field list whose values are compared to determine duplicate events. If a field list is specified, the values of the unique sets of all those fields are used to remove events. For example, if name and deviceCustomNumber1 are specified, and two events contain “Network Usage - Outbound” and “2347896”, only the first event is kept in the search results.

`keepevents` specifies whether to set the fields specified in the field list to NULL or not. When this option is set to True, the values are set to NULL and events are not removed from search results. However, when this option is set to False, duplicate events are removed from the search results. Default: False.

`keepempty` specifies whether to keep events in the search results whose specified fields contain NULL values. When this option is set to True, events with NULL values are kept, however if this option is set to False, events with NULL values are removed. Default: False.

Example 1

To view events from unique devices:

```
... | dedup deviceAddress
```

Example 2

To view unique deviceEventClassId events from unique devices:

```
... | dedup deviceEventClassId deviceAddress
```

Example 3

To view the className in events with Java exceptions in the message field:

```
exception | <rex_expression> | dedup 5 className
```

In the above example, rex expression is not shown in detail however this expression extracts the class name in a field called className, which the dedup operator acts upon.

eval

Displays events that match the resultant of the specified expression. The expression can be a mathematical, string, or Boolean operation and is evaluated when the query is run. The resulting value of the expression is assigned to a field name (as specified in the expression). Once a new field has been defined by the eval operator in a query, this field can be used in the query for further refining the search results (see Example #3 below, in

which a new field "Plus" is defined by the eval operator; this field is then used by the sort operator.)

Usage

```
... | eval <expression>
```

<expression> is a mathematical, string, or Boolean operation; for example, `total_bytes=bytesIn + bytesOut`.

Notes

Typically, a `cef` or `rex` operator (to extract fields from matching events) precedes the `eval` operator, as shown in the examples below. However, you can use the `eval` operator on a field that has been defined by a previous `eval` operator in a query.

Example 1

If the Category Behavior is "Communicate", then assign the value "communicate" to a new field "cat"; otherwise, assign the value "notCommunicate" to it.

```
_storageGroup IN ["Default Storage Group"] | cef categoryBehavior |
eval cat=if(categoryBehavior== "/Communicate", "communicate",
"notCommunicate")
```

Example 2

Append the word, "END", at the end of extracted event name. For example, if event name is "Logger Internal Event", after the `eval` operation it is "Logger Internal EventEND" and is assigned to a new field, "fullname".

```
logger | cef msg name | eval fullname=name + "END"
```

Example 3

Add 100 to the value of `bytesIn` and assign it to a new field, "Plus". Then, sort the values assigned to "Plus" in ascending order.

```
_storageGroup IN ["Default Storage Group"] | cef bytesIn bytesOut
name | eval Plus=bytesIn +100 | sort Plus
```

extract

Extracts key value pairs from raw events.

Usage

```
... | extract [pairedlim="<delimiters>"] [kvdelim="<delimiters>"]
[maxchars=<n>] fields="key1,key2,key3..."
```

`pairedlim` is a delimiter (or a list of delimiters) that separates one key-value pair from another key-value pair in an event. By default, semi colon, pipe, and comma (`;` `|` `,`) are used.

`kvdelim` is a delimiter (or a list of delimiters) that separates a key from its value. By default, `"="`.

`maxchars` is the maximum number of characters in an event that would be scanned for extracting key value pairs. By default, 10240.

`fields` is a key (or a list of comma-separated keys) whose values you want to display in the search results. For example, if you want to display the Name Age, and Location values from this event:

Name:Jane | Age:30 | Location:LA

Then, extract the "Name", "Age", and "Location" keys and list them in the `fields` list.

Understanding how the operator works:

The key represents a field in the raw event and its value consists of the characters that appear after the key until the next key in the event. The following raw event is used to illustrate the concept:

```
[Thu Jul 30 01:20:06 2009] [error] [client 69.63.180.245] PHP Warning:  
memcache_pconnect() [<a  
href='function.memcache-pconnect'>function.memcache-pconnect</a>]: Can't connect to  
10.4.31.4:11211
```

To extract the URL from the above event, you can define these key-pair delimiters, which separate the key-value pairs in the event:

- Greater than sign (" $>$ ")
- Square bracket ("[")

And, define this key delimiter, which separates the key from its value:

- Equal to sign ("=")

Thus, the following command will extract the URL

```
... | extract pairdelim= ">\" kvdelim= "=" fields="<a href"
```

The key value pairs in the event will be: [

The key in the event will be: <a href

The extracted URL will be: 'function.memcache-pconnect'

Notes

This operator only works on raw events. That is, you cannot extract key value pairs from CEF events or the fields defined by the `rex` operator.

You can specify the `pairdelim` and `kvdelim` delimiters in the `extract` operator command to extract keys and their values. However, if you want to determine the key names that these delimiters will generate, use the `keys` operator as described in ["keys" on page 540](#). The `keys` operator can only be used to determine keys; you cannot pipe those keys in the `extract` operator. That is, `...| keys | extract fields=field1` is incorrect.

The keys specified in the fields list can be used further in the pipeline operations. For example, `...| extract pairdelim= "|" kvdelim= ":" fields= "count" | top count`

If none of the specified `pairdelim` characters exist in an event, the event is not parsed into key value pairs. The whole event is skipped. Similarly, if the specified `kvdelim` does not exist, values are not separated from the keys.

To specify double quotes (") as the delimiter, enter it within the pair of double quotes with backslash(\) as the escape character. For example, "=\"|. Similarly, use two backslashes to treat a backslash character literally. For example, "\\\".

Example

```
... | extract pairdelim= "|" kvdelim= ":" fields=
      "Name, Age, Location"
```

Extracts values from events in this format:

Name:Jane | Age:30 | Location:LA

fields

Includes or excludes specified fields from search results.

Usage

```
... | fields [(+ | -)] <field>)+
```

+ includes only the specified field or fields in the search results. This is the default.

- excludes only the specified field or fields from the search results.

Notes

Typically, the <field> list contains event fields available in the Logger schema or user-defined fields created using the `rex` operator prior in the query, as shown in the examples below. However, fields might also be defined by other operators such as the `eval` operator.

The + and - can be used in the same expression when multiple fields are specified. For example, `| fields + name - agentType`

A complete field name must be specified for this operator; wildcard characters in a field name are not supported.

When this operator is included in a query, select **User Defined Fieldsets** from the System Fieldsets list to view the search results.

Example 1

```
... | fields - agentType + categorySignificance
```

Example 2

```
... | fields - name
```

head

Displays the first <N> lines of the search results.

Usage

```
... | head [<N>]
```

<N> is the number of lines to display. Default: 10, if <N> is not specified.

Notes

When this operator is included in a query, the search results are not previewable. That is, the query must finish running before search results are displayed.

Example

```
... | head
```

keys

Identifies keys in raw events based on the specified delimiters.

Usage

```
... | keys [pairedelim= "<delimiters>"] [kvdelim= "<delimiters>"]  
[limit=<n>]
```

`pairedelim` is a delimiter (or a list of delimiters) that separates one key-value pair from another key-value pair in an event. By default, semi colon, pipe, and comma (; | ,) are used.

`kvdelim` is a delimiter (or a list of delimiters) that separates a key from its value. By default, "=".

`limit` is the maximum number of key value pairs to find. There is no default or maximum number for this parameter.

Notes

This operator only works on raw events. That is, you cannot identify key value pairs from CEF events or fields defined by the `rex` operator.

Although this operator is not required to determine keys, it is recommended that you use it to first determine the keys whose values you want to obtain using the `extract` operator. This operator returns aggregated results. Therefore, the search results list the keys found in the matching events and their counts.

The `keys` operator can only be used to determine keys; you cannot pipe those keys in the `extract` operator. That is, `| keys | extract fields=field1` is incorrect.

If a key value is blank (or null), it is ignored and not counted toward the number of hits.

For example, for the following event data:

```
Date=3/24/2011 | Drink=Lemonade  
Date=3/23/2011 | Drink=  
Date=3/22/2011 | Drink=Coffee
```

Search Query: `keys pairedelim= "|" kvdelim= "="`

Search Result: Date, 3 hits and Drink, 2 hits

If none of the specified `pairedelim` characters exist in an event, the event is not parsed into key value pairs. The whole event is skipped. Similarly, if the specified `kvdelim` does not exist, values are not separated from the keys.

To specify double quotes (") as the delimiter, enter it within the pair of double quotes with backslash(\) as the escape character. For example, `"=\"|"`. Similarly, use two backslashes to treat a backslash character literally. For example, `"\\"`.

Example 1

```
... | keys pairdelim= "|" kvdelim= "="
```

Identifies keys (Date and Drink) in event of this format:

Date=3/24/2011 | Drink=Lemonade.

Example 2

```
... | keys pairdelim= "," kvdelim= ">="
```

Identifies keys (Path and IPAddress) in the event of this format:

Path>c:\usr\log, IPAddress=1.1.1.1

parse

Applies the named parser to the matching events of a search query. The parser definition for the specified parser name must exist before it can be used in a query.

Usage

```
... | parse <parser_name>
```

<parser_name> is the name of the parser to use.

The `parse` operator is useful in parsing the non-CEF (unstructured textual) data stored on Logger and parsing it into specific fields according to the parser's definition.

Once parsed into fields, this data can be used further in search operations. For example, the following `parse` operator parses the events using a user-defined parser "Web Server Access Logs" such that "username", "login_status", "num_attempts" fields are created. You can use these created fields further in a pipeline query to display the top 10 user names that resulted in the maximum failed login attempts and the number of attempts they made.

```
... | parse Web Server Access Log | where login_status = "failed" |
top username num_attempts
```

Because the parser definitions are `rex` or `extract` expressions, they create additional fields to contain values that match the specified expression. These fields are displayed in the Search Results just like the results of any `rex` or `extract` expression. Therefore, in the above example, three additional fields will be added to the Search Results—username, login_status, num_attempts.

An additional field called "parser" is also added to the Search Results when the `parse` operator is used in a search query.

This field contains the name of the parser when the parser is able to parse one or more fields specified in the definition for the matching events. If the event was not parsed successfully, if no parser is defined for the source type, or if there is no source type, this field displays, this field contains "Not parsed". Similarly, the field contains the value "not parsed" when the parser definition is not able to parse any fields of the matching event.

You can also use this field to find out events that were successfully parsed or did not parse, as shown in the following example:

```
... | parse Apache Access Log | where parser = "not parsed"
```

When to use the parse operator: When non-CEF events are received through TCP or UDP receivers on Logger, they are not associated with a source type and thus a parser

definition. Therefore, such events not parsed automatically. Similarly, non-CEF events stored on Logger version 5.2 or earlier are not parsed as the parser feature did not exist in those versions. If you need such events parsed when they match a query, use the `parse` operator.

When an event for which a defined source type exists on Logger is parsed through the `parse` operator, it can result in the creation of multiple user-defined fields—through the parser associated with the source type and through the parser you specified in the `parser pipeline` command. If both parsers create unique field names, all those fields are created when a query that matches the event is run. If the parsers specify one or more same name fields, the field names specified in the `parse` operator parser take precedence as this parser is applied last.

Example:

```
... | parse Web Server Access Log | where url CONTAINS ".org" | top url
```

rare

Lists the search results in a tabular form of the least common values for the specified field. That is, the values are listed from the lowest count value to the highest.

When multiple fields are specified, the count of unique sets of all those fields is listed from the lowest to highest count.

Usage

```
... | rare <field1> <field2> <field3> ...
```

Notes

Typically, the `<field>` list contains event fields available in the Logger schema or user-defined fields created using the `rex` or `eval` operators prior in the query, as shown in the examples below. However, fields might also be defined by other operators such as the `eval` operator.

A chart of the search results is automatically generated when this operator is included in a query. You can click on a charted value to quickly filter down to events with specific field values. For more information, see [“Chart Drill Down” on page 114](#).

If multiple fields are specified, separate the field names with a white space or a comma.

Example

```
... | rare deviceEventCategory
```

regex

Selects events that match the specified regular expression.

Usage

```
... | regex <regular_expression>
```

OR

```
... | regex <field> (=|!=) <regular_expression>
```

Notes

Regular expression pattern matching is case insensitive.

The first usage (without a field name) is applied to the raw event. While the second usage (with a field name), is applied to a specific field.

If you use the second usage (as shown above and in the Example #2 below), either specify an event field that is available in the Logger schema or a user-defined field created using the `rex` or `eval` operators.

Example 1

```
... | regex "failure"
```

Example 2

```
... | regex deviceEventCategory != "fan"
```

rename

Renames the specified field name.

Usage

```
...| rename <field> as <new_name>
```

<field> is the name of an event field that is available in the Logger schema or a user-defined field created using the `rex` or `eval` operator.

<new_name> is the new name you want to assign to the field.

Notes

An additional column is added to the search results for each renamed field. The field with the original name continues to be displayed in the search results in addition to the renamed field. For example, if you rename `deviceEventCategory` to `Category`, two columns are displayed in the search results: `deviceEventCategory` and `Category`.

You can include the wildcard character, `*`, in a field name. However, you must enclose the field that contains a wildcard character in double quotes (`" "`). For example:

```
...| rename "*IPAddress" as "*Address"
```

OR

```
...| rename "*IPAddress" as Address
```

If a field name includes a special character (such as `_`, a space, `#`, and so on), it should be included in double quotes (`" "`) in the `rename` operator expression. For example:

```
...| rename src_ip as "Source IP Address"
```

If the resulting field of a `rename` operation includes a special character, it must be enclosed in double quotes (`" "`) whenever you use it in the pipeline operator expression. For example,

```
...| rename src_ip as "Source IP Address" | top "Source IP Address"
```

The internal field names (that start with `"_raw"`) cannot be renamed.

The renamed fields are valid only for the duration of the query.

The resulting field of a `rename` operation is case sensitive. When using such a field in a search operation, make sure that you use the same case that was used to define the field.

When you export the search results of a search query that contains the `rename` expression, the resulting file contains the renamed fields.

Example 1

```
...| rename src_ip as IPAddress
```

Example 2

```
...| rename src_ip as "Source IP Address"
```

replace

Replaces the specified string in the specified fields with the specified new string.

Usage

```
<orig_str> with <new_str> [in <field_list>]
```

`<orig_str>` is the original string you want to replace. (See Notes for more details.)

`<new_str>` is the new string you want to replace with. (See Notes for more details.)

`<field_list>` is the optional, however highly recommended. See Notes for details.

Notes

Even though the field list is optional for this command, HP strongly recommends that you specify the fields on which the `replace` operator should act in this command.

If you skip the field list, the `replace` operator acts on the fields that have been either explicitly defined using the `cef`, `rex`, and `eval` operators preceding the `replace` command, or any fields that were used in other operator commands that preceded the `replace` operator command. For example, the `replace` command acts on `deviceEventCategory` in all of the following cases and replaces all instances of "EPS" with "Events":

```
...| replace *EPS* with *Events* in deviceEventCategory
...| cef deviceEventCategory | replace *EPS* with *Events*
...| top deviceEventCategory | replace *EPS* with *Events*
```

An additional column of the same name is added to the search results for each field in which string is replaced. The column with the original value continues to be displayed in the search results in addition to the column with replaced values. For example, if you replace "err" with "Error" in the "message" column, an additional "message" column is added to the search results that contains the modified value.

If you want to replace the entire string, specify it in full (as it appears in the event). For example, "192.168.35.3".

If you want to replace a part of the string, include wildcard character (*) for the part that is not going to change.

For example, if the original string (the string you want to replace) is "192.168*", only the 192.168 part in an event is replaced. The remaining string is preserved. As a result, if an event contains 192.168.35.3, only the first two bytes are replaced. The rest (35.3) will be

preserved. Similarly, if the event contains 192.168.DestIP, DestIP will be preserved. However, if the event contains the string 192.168, it will not be replaced.

If both, the original and the new strings contain wildcard characters, the number of wildcard characters in the *original* string must match the number of wildcard characters in the *new* string.

```
... | replace "*.168.*" with "*.XXX.*"
```

If the original or the new string includes a special character such as / or ?, enclose the string in double quotes (" "):

```
... | replace "/Monitor" with Error
```

You can replace multiple values for multiple fields in a single operation by separating each expression with a comma (.). Note that you must specify the field list after specifying the "with" expression for all values you want to replace, as shown in the following example:

```
... | replace "Arc*" with HP, "cpu:100" with EPS in deviceVendor,
deviceEventClassId
```

The original string is case-insensitive. Therefore, the string "err" will replace an event that contains "Err".

Example 1

Replace any occurrence of "a" with "b" but the characters preceding "a" and succeeding it are preserved.

```
... | replace *a* with *b*
```

Example 2

Replace any occurrence of "a" with "b" without retaining any characters preceding or succeeding "a".

```
... | replace *a* with b in name
```

rex

Extracts (or capture) a value based on the specified regular expression or extract and substitute a value based on the specified "sed" expression. The value can be from a previously specified field in the query or a raw event message.

Usage

```
... | rex <regular_expression containing a field name>
```

OR

```
... | rex field = <field> mode=sed "s/<string to be
substituted>/<substitution value>"
```

Understanding how extraction works:

When the value is extracted based on a regular expression, the extracted value is assigned to a field name, which is specified as part of the regular expression. The syntax for defining the field name is **?<fieldname>**, where *fieldname* is a string of alphanumeric characters. Using an underscore ("_") is not recommended.

We use the following event to illustrate the power of rex.

[Thu Jul 30 01:20:06 2009] [error] [client 69.63.180.245] PHP Warning: Can't connect to 10.4.31.4:11211

If you want to extract any IP address from the above event and assign it to a field called "IP_Address", you can simply specify the following rex expression:

```
| rex "(?<IPAddress>\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```

However, if you wanted to extract the IP address after the word "client" from the following event and assign it to a field called "SourceIP", you will need to specify a start and end point for IP address extraction so that the second IP address in the event is not captured. The starting point in this event can be "[client" and the end point can be "]". Thus, the rex expression will be:

```
| rex "\[client (?<SourceIP>[^\]]*)"
```

In this rex expression **?<SourceIP>** is the field name defined to capture IP address and "client " specifies the text or point in the event AFTER which data will be extracted. The `[^\]]*` expression will match every character that is not a closing right bracket, therefore, for our example event, the expression will match until the end of the first IP address and not the second IP address that appears after the word "to".

Understanding how substitution works:

When the `rex` operator is used in `sed` mode, you can substitute the values of extracted fields with the values you specify. For example, if you are generating a report of events that contain credit card numbers, you might want to substitute the credit card numbers to obfuscate the real numbers.

The substitution only occurs in the search results. The actual event is not changed.

In the following example, the credit card numbers in the CCN field are substituted with "xxxx", thus obfuscating sensitive data:

```
| rex field=CCN mode=sed "s/*/XXXX/g"
```

The `/g` at the end of the command indicates a global replace, that is, all occurrences of the specified pattern will be replaced in all matching events. If `/g` is omitted, only the first occurrence of the specified pattern in each event is replaced.

Multiple substitutions can be performed in a single command, as shown in the following example. In this example, the word "Authentication" is substituted with "xxxx" globally (for all matching events), the first byte of the agent address that start with "192" is substituted with "xxxx" and an IP address that starts with "10" is substituted with "xxxx".

```
| rex field=msg mode=sed "s/Authentication/xxxx/g" | rex  
field=agentAddress mode=sed "s/192/xxxx/g" | rex field=dst mode=sed  
"s/10./xxxx/g"
```

Notes

A detailed tutorial on the `rex` operator is available at [Appendix B, Using the Rex Operator, on page 553](#).

A Regex Helper tool is available for formulating regular expressions of fields in which you are interested. The Regex Helper parses an event into fields. Then, you select the fields that you want to include in the rex expression. The regular expression for those fields is automatically inserted in the Search box. For detailed information on the Regex Helper tool, see ["Regex Helper Tool" on page 101](#).

The extracted values are displayed as additional columns in the All Fields view (of the System Fieldsets). To view only the extracted columns, select **User Defined Fieldsets** from the System Fieldsets list. In the above example, an additional column with heading "SourceIP" is added to the All Fields view; IP address values extracted from events are listed in this column.

If you want to use other search operators such as fields, sort, chart, and so on to refine your search results, you must first use this operator to extract those fields.

Example 1

The following example extracts name and social security number from an event that contains data in name:John ssn:123-45-6789 format and assigns them to Name and SSN fields:

```
... | rex "name: (?<Name>.*) ssn: (?<SSN>.*)"
```

Example 2

The following example extracts URLs from events and displays the top 10 of the extracted URLs:

```
... | rex "http://(?<URL>[^\ ]*)" | top URL
```

Example 3

The following example substitutes the last four digits of social security numbers extracted in the first event with XXXX:

```
... | rex field=SSN mode=sed "s/-\d{4}/-XXXX/g"
```

sort

Sorts search results as specified by the sort criteria.

Usage

```
... | sort [<N>] ((+ | -) field)+
```

+ Sort the results by specified fields in ascending order. This is the default.

- Sort the results by specified fields in descending order.

<N> Keep the top N results, where N can be a number between 1 and 10,000. Default: 10,000.

Notes

Typically, the <field> list contains event fields available in the Logger schema or user-defined fields created using the `rex` operator prior in the query, as shown in the examples below. However, fields might also be defined by other operators such as the `eval` operator.

Sorting is based on the data type of the specified field.

When multiple fields are specified for a sort operation, the first field is used to sort the data. If there are multiple same values after the first sort, the second field is used to sort within the same values, followed by third field, and so on. For example, in the example below, first the matching events are sorted by "cat" (device event category). If multiple events have the same "cat", those events are further sorted by "eventId".

When multiple fields are specified, you can specify a different sort order for each field. For example, `| sort + deviceEventCategory - eventId`.

If multiple fields are specified, separate the field names with a white space or a comma.

Sorting is case-sensitive. Therefore, "Error:105" will precede "error:105" in the sorted list (when sorted in ascending order).

When a sort operator is included in a query, only the top 10,000 matches are displayed. This is a known limitation and will be addressed in a future Logger release.

When this operator is included in a query, the search results are not previewable. That is, the query must finish running before search results are displayed.

Example

```
... | sort deviceEventCategory eventId
```

tail

Displays the last <N> lines of the search results.

Usage

```
... | tail [<N>]
```

<N> is the number of lines to display. Default: 10, if <N> is not specified.

Notes

When this operator is included in a query, the search results are not previewable. That is, the query must finish running before search results are displayed.

Example

```
... | tail 5
```

top

Lists the search results in a tabular form of the most common values for the specified field. That is, the values are listed from the highest count value to the lowest.

Usage

```
... | top [<n>] <field1> <field2> <field3> ...
```

<n> limits the matches to the top *n* values for the specified fields. Default: 10, if <N> is not specified.

Notes

The fields can be either event fields available in the Logger schema or user-defined fields created using the `rex` or `eval` operators prior in the query. If multiple fields are specified, separate the field names with a white space or a comma.

When multiple fields are specified, the count of unique sets of all those fields is listed from the highest to lowest count.

A chart of the search results is automatically generated when this operator is included in a query. You can click on a charted value to quickly filter down to events with specific field values. For more information, see ["Chart Drill Down" on page 114](#).

To limit the matches to the top n values for the specified fields, specify a value for n . For example, `... | top 5 deviceEventCategory`

Example 1

```
... | top deviceEventCategory
```

Example 2

```
... | top 5 categories
```

transaction

Groups events that have the same values in the specified fields.

Usage

```
... | transaction <field1> <field2>... [maxevents=<number>]
[maxspan=<number>[s|m|h|d]] [maxpause=<number>[s|m|h|d]]
[startswith=<reg_exp>] [endswith=<reg_exp>]
```

`field1`, `field2` is a field or a comma-separated field list whose values are compared to determine events to group. If a field list is specified, the values of the unique sets of all those fields are used to determine events to group. For example, if `host` and `portNum` are specified, and two events contain "hostA" and "8080", the events are grouped in a transaction.

`maxevents` specifies the maximum number of events that can be part of a single transaction. For example, if you specify 5, after 5 matching events have been found, additional events are not included in the transaction. Default: 1000

`maxspan` specifies the limit on the duration of the transaction. That is, the difference in time between the first event and all other events in a transaction will never be more than the specified `maxspan` limit. For example, if you specify `maxspan=30s`, the event time of all events within the transaction will be at most 30 seconds more than the event time of the first event in the transaction. Default: Unlimited

`maxpause` specifies the length of time by which consecutive events in a transaction can be apart. That is, this option ensures that events in a single transaction are never more than the `maxpause` value from the previous event in the transaction. Default: Unlimited

`startswith` specifies a regular expression that is used to recognize the beginning of a transaction. For example, if a transaction operator includes `startswith= "user [L|l]login"`, all events are scanned for this regular expression. When an event matches the regular expression, a transaction is created, and subsequent events with matching fields are added to the transaction.



Note

The regular expression is applied to the raw event, not to a field in an event.

`endswith` specifies a regular expression that is used to recognize the end of an existing transaction. That is, an existing transaction is completed when an event matches the specified "endswith" regular expression. For example, if a transaction operator includes

endswith= "[L|l]ogout", any event being added to a transaction is checked, and if the regular expression matches the event, the transaction is completed.



Note

The regular expression is applied to the raw event, not to a field in an event.

Notes

Several of the above options specify "conditions to end" a transaction. Therefore, when multiple "end conditions" are specified in a transaction operator, the first end condition that occurs will end the transaction even if the other conditions have not been satisfied yet. For example, if `maxspan` is reached but `maxevents` has not been reached, or if the `endswith` regular expression is matched but `maxevents` has not been reached.

Understanding how the transaction operator works:

A transaction is a set of events that contain the same values in the specified fields. The events may be further filtered based on the options described above, such as `maxspan`, `maxpause`, and so on. In addition to grouping events, the transaction operator adds these fields to each event: `transactionid`, `duration`, and `eventcount`. These fields are displayed in the Search Results as separate columns.

A `transactionid` is assigned to each transaction when the transaction completes. Transaction IDs are integers, assigned starting from 1 for the transactions (set of events) found in the current query. All events in the same transaction will have the same transaction ID.

If an event does not belong to any transaction found in the current query, it is assigned the transaction ID 0. For example, in a `transaction` operator with a `startswith` regular expression, if the first event in the pipeline does not match the regular expression, that event is not part of the transaction, and is assigned transaction ID 0.

The `duration` is the time in milliseconds of the duration of a transaction, which is the difference between the event time of the last event in the transaction and the first event in the transaction. The `duration` field for all events in a transaction is set to the `duration` value of the transaction.

The `eventcount` displays the number of events in a transaction.

Example 1

To view source addresses accessed within a 5-minute duration:

```
... | transaction sourceAddress maxspan=5m
```

Example 2

To group source addresses by source ports and view 5 events per group:

```
...| transaction sourceAddress sourcePort maxevents=5
```

Example 3

To group users and URLs they accessed within a 10-minute duration:

```
... | transaction username startswith= "http://" maxspan=10m
```

Example 4

To view login transactions from the same session ID and source address in a 1-hour duration:

```
... | transaction sessionId sourceAddress maxspan=1h startswith=
"user [L|l]login"
```

where

Displays events that match the criteria specified in the "where" expression.

Usage

```
... | where <expression>
```

<expression> can be any valid field-based query expression, as described in ["Field-based Search" on page 80](#).

Notes

<expression> can only be a valid field-based query expression. Arithmetic expressions or functions are not supported.

Example 1

```
... | where eventId is NULL
```

Example 2

```
... | where eventId=10006093313 OR deviceVersion CONTAINS
"4.0.6.4924.1"
```

Example 3

```
... | where eventId >=10005985569 OR categories= "/Agent/Started"
```


Appendix B

Using the Rex Operator

The `rex` operator is a powerful operator that enables you to extract information that matches a specified regular expression and assigns it to a field, whose field name you specify. You can also specify an optional start point and an end point in the rex expression between which the information matching the regular expression is searched.

This appendix describes the `rex` search operator in detail. It includes information on the following topics.

[“Syntax of the rex Operator” on page 553](#)

[“Ways to Create a rex Expression” on page 554](#)

[“Example rex Expressions” on page 555](#)

When a rex expression is included in a search query, it must be preceded by a basic search query that finds events from which the rex expression will extract information. For example:

```
failed | rex "(?<srcip>[^\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```

Syntax of the rex Operator

```
| rex "text1(?<field1>text2regex)"
```

text1—The text or point in the event AFTER which information extraction begins. The default is the beginning of the event.

text2—The text or point in the event at which information extraction ends.

field1—The name of the field to which the extracted information is assigned.

regex—The pattern (regular expression) used for matching information to be extracted between *text1* and *text2*.



If you are an experienced regular expression user, see the Note in the next section for a quick understanding of how rex enables you to capture named input and reference it for further processing.

Understanding the rex Operator Syntax

Extract all information AFTER *text1* and until *text2* that matches the specified *regex* (regular expression) and assign TO *field1*.

- **text1** and **[text2]** can be any points in an event—start and end of an event, specific string in an event (even if the string is in the middle of a word in the event), a specific number of characters from the start or end of an event, or a pattern.
- To specify the next space in the event as **text2**, enter **[^]**.

This is interpreted as “not space.” Therefore, entering a “not” results in the capture to stop at the point where the specified character, in this case, a space, is found in the event.

- To specify **[text2]** to be the end of the line, enter **[^\$]**.

This is interpreted as “not end of line.” Therefore, when an end-of-line in an event is encountered, the capture will stop at that point. The **[^\$]** usage only captures one character if it is not an end-of-line character. However, by specifying **[^\$]*** in a rex expression, the usage captures all characters until end-of-line.

You can also specify **.*** to capture all characters in an event instead of **[^\$]**. Examples in this document, however, use **[^\$]**.

- Any extra spaces within the double quotes of the rex expression are treated literally.
- The characters that need to be escaped for rex expressions are the same as the ones for regular expressions. Refer to a regular expressions document of your choice to obtain a complete list of such characters.
- Information captured by a rex expression can be used for further processing in a subsequent rex expression as illustrated in the following example in which an IP address is captured by the first rex expression and the network ID (assuming the first three bytes of the IP address represent it) to which the IP address belongs is extracted from the captured IP address:

```
logger | rex "(?<srcip>[^ ]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
| rex field=srcip "(?<netid>\d{1,3}\.\d{1,3}\.\d{1,3})"
```



Note

If you are an experienced regular expression user, you can interpret the rex expression syntax as follows:

```
rex "(?<field1>regex)"
```

where the entire expression in the parentheses specifies a named capture. That is, the captured group is assigned a name, which can be referenced later for further processing. For example, in the following expression “srcip” is the name assigned to the capture.

```
failed | rex "(?<srcip>[^ ]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```

Once named, use “srcip” for further processing as follows:

```
failed | rex "(?<srcip>[^ ]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" |
top srcip
```

Ways to Create a rex Expression

You can create a rex expression in two ways:

- **Manually**—Follow the syntax and guidelines described in this appendix to create a rex expression to suit your needs.
- **Regex Helper**—Use the Regex Helper tool, as described in [“Regex Helper Tool” on page 101](#). This tool not only simplifies the process, it also makes it less error prone and more efficient.

Creating a rex Expression Manually

Start with a simple search that finds the events that contains the information in which you are interested. Once the events are displayed, identify a common starting point in those events that precedes the information.

For example, you are interested in extracting the client IP address, which always appears after the word "[client" in the following event.

```
[Thu Jul 30 01:20:06 2009] [error] [client 69.63.180.245] PHP
Warning: memcache_pconnect() [<a
href='function.memcache-pconnect'>function.memcache-pconnect</a>]:
Can't connect to 10.4.31.4:11211
```

Therefore, "[client" is the starting point. A good end point is the "]" after the last byte of the client IP address. Now, we need to define the regular expression that will extract the IP address. Because in this example, only the client IP address appears after the word "client", we use "*" as the regular expression, which means "extract everything". (We could be more specific and use `\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}` for the IP address.) We assign the extracted IP address to a field name "clientIP". We are almost ready to create a rex expression, except that we need to escape the "[" and "]" characters in the expression. The escape character to use is "\".

Now, we are ready to create the rex expression to extract the IP address that appears after the word "client" in the event shown above.

```
| rex "\[client(?:<clientip>[^\]]*)" "
```

Example rex Expressions

This section contains several sample examples for extracting different types of information from an event. The specificity of the information extracted increases with each example. Use these examples as a starting point for creating rex expressions to suit your needs. Also, use the Regex Helper tool that simplifies rex expression creation.

This event is used as an example to illustrate the information the following rex expressions will extract:

```
2010/07/01 13:46:00 PDT    unknown    Local    ArcSight    Logger    4.5.0.4836.0    eps:100
CEF:0|ArcSight|Logger|4.5.0.4836.0|eps:100|Logger Internal Event|1| cat=/Monitor/Receiver/All/EP5 cs2=SinceLastMonitorEvent cnt=1 dvc=192.168.36.3 cs1=Events/sr
2010/07/01 13:46:00 PDT    unknown    Local    ArcSight    Logger    4.5.0.4836.0    eps:100
CEF:0|ArcSight|Logger|4.5.0.4836.0|eps:100|Logger Internal Event|1| cat=/Monitor/Forwarder/All/EP5 cs2=SinceLastMonitorEvent cnt=1 dvc=192.168.36.3 cs1=Events/sr
```

- Capture matching events from the left of the pipeline and assign them to the field, message. The entire event is assigned to the "message" field.

```
| rex "(?<message>[^\$]*)"
```

This expression extracts the entire event (as shown above), starting at the word "CEF:0".

- Specifying the starting point as number of characters from the start of an event instead of a specific character or word

```
| rex "[a-zA-Z0-9:\.\s]{16}(?<message>[^\$]*)"
```

This expression starts extracting after 16 **consecutive** occurrences of the characters specified for *text1*—alphanumeric characters, colons, periods, or spaces. Although the first 16 characters of the first event are "CEF:0|ArcSight|L", the extraction does not

begin at “ogger|4.5.0...” because the pipeline character is not part of the characters we are matching, but this character is part of the beginning of the event. Therefore, the first 16 consecutive occurrences are “Logger Internal “. As a result, information starting at the word “Event”, is extracted from our example event.

- Extract a specified number of characters instead of specifying an end point such as the next space or the end of the line

```
| rex "[a-zA-Z0-9:\.\s]{16} (?<message>[^$]{5})"
```

This expression only extracts the word “Event”. (See the previous sample rex expression for a detailed explanation of the reason extraction begins at the word “Event”.)

- Extract everything after “CEF:0|” into a field, `message`. Then, pipe events for which the `message` field is not null through another rex expression to extract the IP address contained in the matching events and assign the IP addresses to another field, `msgip`. Only display events where `msgip` is not null.

```
| rex "CEF:0|(?<message>[^\$]*" | where message is not null |
rex "dvc=(?<msgip>[ ^]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" |
where msgip is not null
```



The “:” and “=” characters do not need to be escaped; however, “|” must be escaped. The characters that need to be escaped for rex expressions are the same as the ones for regular expressions. Refer to a regular expressions document of your choice to obtain a complete list of such characters.

This expression extracts the device IP address from the event.

The following `rex` examples use this event for illustration:

Nov 10 03:04:24 192.168.20.114 192.168.20.113 [REDACTED]C007:4D28:EviPackets:Line 16: 'New Group', 'My [REDACTED]150', '11/10/2005 11:02:05.000', '21561', '11/10/2005 11:02:05.000', '3106004', 'generator', '1', '192.168.20.111', 'http:80', '192.168.20.112', '32771', 'tcp', 'Alert', '47302', '47285', 'RPC Incomplete Segment', '0', '0', '00:00:00:00:00:00:00:00', '00:00:00:00:00:00:00:00'

- Extract the first two IP addresses from an event and assign them to two different fields, IP1 and IP2.

```
| rex "(?<IP1>[^\$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" | rex
"\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}(?<IP2>[^\$]\d{1,3}\.\d{1,3}\.
\d{1,3}\.\d{1,3})"
```

This expression extracts the first and second IP addresses in the above event.

Because the two IP addresses are right after one another in this event, you can also specify the extraction of the two IP addresses in a single rex expression as follows:

```
| rex
"(?[^$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})(?[^$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})"
```



Do not specify a space in the above expression.

- Building on the previous example, add a new field called Ignore. Assign the value "Y" to this field if the two IP addresses extracted in the previous example are the same and assign the value "N" if the two IP addresses are different. Then, list the top IP1 and IP2 combinations for events for which Ignore field is "N".


```
| rex (?<IP1>[^$]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" | rex
"\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3}(?<IP2>[^$]\d{1,3}\.\d{1,3}\.
\d{1,3}\.\d{1,3})" | eval Ignore=if(IP1==IP2,"Y","N") | where
Ignore="N" | top IP1 IP2
```



The eval command uses double == to equate the two fields.

- Information captured by a rex expression can be used for further processing in a subsequent rex expression as illustrated in the following example. The first IP address is captured by the first rex expression and the network ID (assuming the first three bytes of the IP address represent it) to which the IP address belongs is extracted from the captured IP address:

```
logger | rex "(?<srcip>[^ ]\d{1,3}\.\d{1,3}\.\d{1,3}\.\d{1,3})" | rex field=srcip
"(?<netid>\d{1,3}\.\d{1,3}\.\d{1,3})"
```

The following rex example uses this event for illustration:

```
127.0.0.1 - name [10/Oct/2010:13:55:36 -0700] "GET /apache_pb.gif HTTP/1.0" 200 2326
"http://www.example.com/start.html" "Mozilla/4.08 [en] (Win98; I ;Nav)" ¶
```

- Extract all URLs from events and generate a chart of the URL counts, excluding blank URLs

```
| rex "http://(?<customURL>[ ^ ]*)" | where customURL is not null
| chart count by customURL | sort - customURL
```



- The events contain the URL string in "http://" format.
- Meta character / needs to be enclosed in square brackets [] to be treated literally.

The following rex example uses this event for illustration:

1	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local	root
RAW	Feb 25 14:03:24 beach login(pam_unix)[123]: session closed for ssaf root	
2	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local	sysadmin
RAW	Feb 25 14:03:24 beach sshd(pam_unix)[123]: authentication failure; logname=uid=0 euid=0 tty=ssh ruser= rhost=123.123.123.123 ssaf =sysadmin	
3	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local	piadmin
RAW	Feb 25 14:03:24 beach sshd(pam_unix)[123]: session closed for ssaf piadmin	
4	2010/04/06 22:11:46 CDT 10.0.10.222 [Raw_Syslog] Local	sysadmin
RAW	Feb 25 14:03:24 beach sshd(pam_unix)[123]: session opened for ssaf sysadmin by (uid=500)	

- Extract the first word after the word "user" (one space after the word) or "user=". The word "user" is case-insensitive in this case and must be preceded by a space character. That is, words such as "ruser" and "suser" should not be matched.

```
| rex "\s[u|U][s|S][e|E][r|R][\s|=](?<CustomUser>[ ^ ]*)"
```


Appendix C

Logger Audit Events

This appendix describes Logger's audit events in detail. It includes information on the following topics:

- “Types of Audit Events” on page 559
- “Information in an Audit Event” on page 560
- “Platform Events” on page 560
- “Logger Application Events” on page 566

You can forward the Logger audit events, which are in Common Event Format (CEF), to ArcSight ESM directly for analysis and correlation. Use the Audit Forwarding feature (as described in “Audit Forwarding” on page 371) to forward the events.

For more information about the Common Event Format (CEF), refer to Implementing ArcSight CEF. To find this guide, search for “ArcSight Common Event Format (CEF) Guide” on the Protect 724 Community at <https://protect724.arcsight.com>.

Types of Audit Events

Two types of audit events are generated on Logger and available for Audit Forwarding to ArcSight ESM.

- “Platform Events” on page 560—related to the Logger hardware/system
- “Logger Application Events” on page 566—related to Logger functions and configuration changes on it

Both types of events are stored in the Logger Internal Storage Group. As a result, these events can be searched using the Logger Search UI. For example, you can search for this platform event:

“/Platform/Authentication/Failure/Password”

In addition to these events, a Logger appliance that has an ArcSight Connector Appliance installed on it generates Connector Appliance audit events. For a list of Connector Appliance audit events, see the Administrator's Guide for Connector Appliance for the version that applies to you.

Information in an Audit Event

A Logger audit event (in CEF format) contains information about the following prefix fields:

- Device Event Class 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 [Regex Query Test] has been added
```

Platform Events

The following table lists the information contained in audit events related to the Logger platform. All events include the following fields.

- duser—UserName
- duid—User ID
- src—IP address of client
- dst—IP address of appliance
- cat—Device Event Category
- cn1—Session number
- cn1label—Session

Additional fields (if applicable) are listed in the following table.

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform: 200	5	/Platform/Authentication/PasswordChange/Failure	Failed password change	
platform: 201	7	/Platform/Authentication/Failure	Failed login attempt	
platform: 202	5	/Platform/Authentication/PasswordChange	Password changed	cs1: Affected User Id cs2: Affected User Login cs3: Affected User Full Name
platform: 203	7	/Platform/Authentication/InactiveUser/Failure	Login attempt by inactive user	
platform: 213	7	/Platform/Configuration/Global/AuditEvents	Audit forwarding modified	cs1: Audit Forwarders
platform: 220	5	/Platform/Certificate/Install	Installed certificate	cs1: Network Protocol

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform:221	7	/Platform/Certificate/Mismatch	Certificate mismatch failure	cs1: Network Protocol
platform:222	1	/Platform/Certificate/Request	Created certificate signing request	cs1: Certificate Signing Request cs2: Network Protocol
platform:224	5	/Platform/Certificate/Regenerate	Re-generate self-signed certificate	cs1: Certificate Signing Request cs2: Network Protocol
platform:226	7	/Platform/Update/Failure/CorruptPackage	Uploaded update file damaged or corrupt	cs1: Error cs2: fname cs3: fsize
platform:227	5	/Platform/Update/Applied	Update installation success	cs1: Update Name cs2: Is Reboot Required
platform:228	7	/Platform/Update/Failure/Installation	Update installation failure	cs1: Error cs2: Update Name
platform:230	3	/Platform/Authentication/Login	Successful login	
platform:234	7	/Platform/Authentication/Failure/LOCKED	Failed login attempt (LOCKED)	
platform:239	3	/Platform/Authentication/Logout	User logout	
platform:240	3	/Platform/Authorization/Groups/Add	Added user group	cn2: Current Number of Users cn3: Current Number of User Rights cs1: Affected Group Name cs2: Affected Group Id flexNumber1: Old Number of Users flexNumber2: Old Number of User Rights
platform:241	3	/Platform/Authorization/Groups/Update	Updated user group	cn2: Current Number of Users cn3: Current Number of User Rights cs1: Affected Group Name cs2: Affected Group Id flexNumber1: Old Number of Users flexNumber2: Old Number of User Rights
platform:242	5	/Platform/Authorization/Groups/Membership/Update/Clear	Removed all members from group	
platform:243	3	/Platform/Authorization/Groups/Membership/Update	Modified user group membership	

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform: 244	3	/Platform/Authorization/Groups/Delete	Deleted user group	cs1: Affected Group Name cs2: Affected Group Id
platform: 245	3	/Platform/Authorization/Users/Add	Added user	cs1: Affected User Id cs2: Affected User Login cs3: Affected User Full Name
platform: 246	3	/Platform/Authorization/Users/Update	Updated user	cs1: Affected User Id cs2: Affected User Login cs3: Affected User Full Name
platform: 247	3	/Platform/Authorization/Users/Delete	Deleted user	cs1: Affected User Id cs2: Affected User Login cs3: Affected User Full Name
platform: 248	3	/Platform/Authentication/Logout/SessionExpiration	Session expired	
platform: 249	7	/Platform/Authentication/AccountLocked	Account locked	
platform: 250	5	/Platform/Storage/RFS/Add	Added remote mount point	cs1: RFS Mount Name cs2: RFS Mount Host and Remote Path
platform: 251	5	/Platform/Storage/RFS/Edit	Edited remote mount point	cs1: RFS Mount Name cs2: RFS Mount Host and Remote Path
platform: 252	7	/Platform/Storage/RFS/Failure	Failed to create remote mount point	cs1: Server cs2: Remote Directory cs3: Mount Name cs4: Mount Type cs5: Username
platform: 253	5	/Platform/Storage/RFS/Remove	Removed remote mount point	cs1: RFS Mount Name cs2: RFS Mount Host and Remote Path
platform: 254	5	/Platform/Storage/SAN/Destroy	Destroyed SAN Logical Unit	cs1: Volume label
platform: 255	5	/Platform/Storage/SAN/Attach	Attached SAN Logical Unit	cn2: Volume size (in MB) cs1: Volume label cs2: World-wide Name cs3: Filesystem type
platform: 256	7	/Platform/Storage/SAN/Detach	Detached SAN Logical Unit	cs1: Storage unit details
platform: 259	5	/Platform/Storage/SAN/Reattach	Reattached SAN Logical Unit	cs1: Volume label cs2: Filesystem type
platform: 260	5	/Platform/Configuration/Network/Route/Update	Static route modified	cs1: Destination cs2: Subnet cs3: Gateway

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform:261	5	/Platform/Configuration/Network/Route/Remove	Static route removed	cs1: Destination cs2: Subnet cs3: Gateway
platform:262	5	/Platform/Configuration/Time	Appliance time modified	cs1: Old Date/Time cs2: New Date/Time cs3: Old Time Zone cs4: New Time Zone
platform:263	5	/Platform/Configuration/Network	NIC settings modified	cs1: NIC cs2: IP Address cs3: Netmask cs4: Speed
platform:264	5	/Platform/Configuration/Network/NTP	NTP server settings modified	cs1: NTP Servers cs2: Is Appliance NTP Server
platform:265	5	/Platform/Configuration/Network/DNS	DNS settings modified	
platform:266	5	/Platform/Configuration/Network/Hosts	Hosts file modified	cs1: Difference from previous hosts file
platform:267	5	/Platform/Configuration/SMTP	SMTP settings modified	cs1: EMail Address cs2: SMTP Server cs3: Backup SMTP Server
platform:268	5	/Platform/Configuration/Network/Route/Add	Static route added	cs1: Destination cs2: Subnet cs3: Gateway
platform:270	5	/Platform/Authorization/Users/Inactive/Disable	Inactive user disabled	cs1: User Login deviceCustomDate1: Date Last Active
platform:280	7	/Appliance/State/Reboot/Initiate	Appliance reboot initiated	
platform:281	3	/Appliance/State/Reboot/Cancel	Appliance reboot canceled	
platform:282	7	/Appliance/State/Shutdown	Appliance poweroff initiated	
platform:284	5	/Platform/Storage/Multipathing/Enable	Enabled SAN Multipathing	cs1: Multipath Configuration
platform:285	5	/Platform/Storage/Multipathing/Disable	Disabled SAN Multipathing	
platform:300	5	/Platform/Certificate/Install	Installed trusted certificate	cs1: Certificate details
platform:301	5	/Platform/Certificate/Revocation/Install	Installed certificate revocation list	cs1: CRL details
platform:302	5	/Platform/Certificate/Delete	Deleted trusted certificate	cs1: Certificate details
platform:303	5	/Platform/Certificate/Revocation/Delete	Deleted certificate revocation list	cs1: CRL details

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform: 304	7	/Platform/Certificate/Install/Failure	Failed installing trusted certificate	cs1: Error cs2: File Size cs3: File Name
platform: 305	7	/Platform/Certificate/Revocation/Install/Failure	Failed installing certificate revocation list	cs1: Error cs2: File Size cs3: File Name
platform: 306	5	/Platform/Process/Start	Start process	cs1: Process Name
platform: 307	5	/Platform/Process/Stop	Stop process	cs1: Process Name
platform: 308	5	/Platform/Process/Restart	Restart process	cs1: Process Name
platform: 310	5	/Platform/Configuration/FIPS/Enable	Enabled FIPS mode	
platform: 311	7	/Platform/Configuration/FIPS/Disable	Disabled FIPS mode	
platform: 312	7	/Platform/Configuration/WebServer/CipherStrength	Web server cipher strength changed	cs1: New Value cs2: Old Value
platform: 320	3	/Appliance/State/Shutdown/Cancel	Appliance poweroff canceled	
platform: 371	5	/Platform/Service/Restart	Restarted OS service	cs1: Service Name
platform: 400	2	/Platform/Diagnostics/Command	Ran diagnostic command	cs1: Diagnostic Command
platform: 407	7	/Platform/Certificate/SSL/Expiration	SSL certificate expiration warning	cs1: Issuer cs2: Subject deviceCustomDate1: Expiration Date
platform: 408	5	/Appliance/State/Startup	Appliance startup completed	deviceCustomDate1: Startup Date
platform: 409	3	/Platform/Configuration/LoginBanner	Configure login warning banner	cs1: Acknowledgement Prompt cs2: Banner Text
platform: 410	5	/Platform/Configuration/Network	Network settings modified	cs1: Gateway cs2: Multi-homing cs3: Hostname
platform: 411	5	/Platform/Authentication/PasswordChange	Automated Password Reset	cn2: User ID cs1: User Login
platform: 412	3	/Platform/Configuration/Locale	Set Locale	cs1: Locale
platform: 440	3	/Platform/Configuration/SNMP	SNMP configuration modified	cn2: Port Number cn3: Refresh Interval cs1: SNMP Enabled cs2: Community String cs3: Listen Address(es)

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform:460	3	/Platform/Network/Aliases/Add	NIC alias added	cs1: NIC cs2: IP Address cs3: Netmask
platform:462	3	/Platform/Network/Aliases/Remove	NIC alias removed	cs1: NIC cs2: IP Address cs3: Netmask
platform:500	5	/Platform/Authorization/Groups/Membership/Remove	Remove member from group	cs1: Affected Group Name cs2: Affected User Login cs3: Affected Group Id cs4: Affected User Id
platform:501	5	/Platform/Authorization/Groups/Membership/Add	Group member added	cs1: Affected Group Name cs2: Affected User Login cs3: Affected Group Id cs4: Affected User Id
platform:502	5	/Platform/Authorization/Users/Groups/Remove	User removed from group	cs1: Affected Group Name cs2: Affected User Login cs3: Affected Group Id cs4: Affected User Id
platform:503	5	/Platform/Authorization/Users/Groups/Add	User added to group	cs1: Affected Group Name cs2: Affected User Login cs3: Affected Group Id cs4: Affected User Id
platform:530	5	/Platform/Configuration/Authentication/Session/Success	Authentication Session settings successfully changed.	cn2: New Value cn3: Old Value cs1: Parameter Changed
platform:540	5	/Platform/Configuration/Authentication/Password/Lockout/Success	Password Lockout settings successfully updated.	cn2: New Value cn3: Old Value cs1: Parameter Changed
platform:550	5	/Platform/Configuration/Authentication/Password/Expiration/Success	Password Expiration settings successfully updated.	cn2: New Value cn3: Old Value cs1: Parameter Changed
platform:560	5	/Platform/Configuration/Authentication/Password/Validation/Success	Password Validation settings successfully updated.	cn2: New Value cn3: Old Value cs1: Parameter Changed
platform:570	5	/Platform/Configuration/Authentication/Password/AutomatedPasswordReset/Success	Password Automated Password Reset setting successfully updated.	cs1: Parameter Changed cs2: New Value cs3: Old Value

Device Event Class ID	Severity	Device Event Category (cat)	Message	Additional Fields
platform: 580	5	/Platform/Configuration /Authentication/Certificate/Success	Client Certificate authentication settings successfully changed.	cs1: Parameter Changed cs2: New Value cs3: Old Value
platform: 590	5	/Platform/Configuration /Authentication/RADIUS/Success	RADIUS authentication settings successfully changed.	cs1: Parameter Changed cs2: New Value cs3: Old Value
platform: 600	5	/Platform/Configuration /Authentication/LDAP/Success	LDAP authentication settings successfully changed.	cs1: Parameter Changed cs2: New Value cs3: Old Value
platform: 610	5	/Platform/Configuration /Authentication/Global/Success	Global Authentication settings successfully changed.	cs1: Parameter Changed cs2: New Value cs3: Old Value

Logger Application Events

The following table lists the information contained in audit events related to various Logger functions and configuration changes on it. The Severity for all Logger application events is 2.

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
Alerts			
logger: 610	/Logger/Component/Alert/Configuration/Add	Alert [name] has been added	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmp HostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:611	/Logger/Component/Alert/Configuration/Delete	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
logger:612	/Logger/Component/Alert/Configuration/Update	Alert [name] has been updated	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmplpAddr dvchost=syslogOrSnmplpHost cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmplpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger:613	/Logger/Component/Alert/Configuration/Enable	Alert [name] has been enabled	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

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger: 614	/Logger/Component/Alert/Configuration/Disable	Alert [name] has been disabled	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmp HostName cn1Label=Syslogor SNMP Destination Port cn1=syslogOrSnmpPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
logger: 615	/Logger/Alert/Configuration/Sent	Alert [name] has been sent	fname=AlertName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=syslogOrSnmpIpAddr dvchost=syslogOrSnmpOr EsmHostName cn1Label=SyslogoOrSnmpOr EsmDestination Port cn1=syslogOrSnmpOrEsmPort cs1Label=Filter cs1=filter cs2Label=Email Destination(s) cs2=emailAddresses
Certificates			
logger: 643	/Logger/Component/Certificate/Configuration/Add	Certificate [name] has been added	fname=alias duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate
logger: 650	/Logger/Component/Certificate/Configuration/Delete	Certificate [name] has been deleted	fname=alias duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate
logger: 651	/Logger/Component/Certificate/Configuration/Update	Certificate [name] has been updated	fname=alias duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Certificate

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
Configuration Backup			
logger: 660	/Logger/Component/ConfigBackup/Configuration/Update	Configuration backup has been updated	fname=Configuration Backup duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Configuration Backup
logger: 661	/Logger/Component/ConfigBackup/Configuration/Enable	Configuration backup has been enabled	fname=Configuration Backup duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Configuration Backup
logger: 662	/Logger/Component/ConfigBackup/Configuration/Disable	Configuration backup has been disabled	fname=Configuration Backup duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Configuration Backup
logger: 665	/Logger/Component/ConfigBackup/Configuration/Backup	Configuration backup succeeded. Transfer process finished.	fname=Configuration Backup fileType=Configuration Backup fpath=pathToBackupFile fsize=fileSizeInByte
ESM Destinations			
logger: 640	/Logger/Component/EsmDestination/Configuration/Add	ESM destination [name] has been added	fname=esmDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=ESM Destination fileId=esmDestinationId dvc=esmDestinationIp dvchost=esmDestinationHost cn1Label=ESM Destination Port cn1=esmDestinationPort cs1Label=ConnectorName cs1=connectorName cs2Label=Connector Location cs2=connectorLocation cs3Label=Logger Location cs3=loggerLocation

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger: 641	/Logger/Component/EsmDestination/Configuration/Delete	ESM destination [name] has been deleted	fname=esmDestinationName duser=UserName duid=userId cs4=sessionId file cs4Label=Session ID fileType=ESM Destination fileId=esmDestinationId dvc=esmDestinationIp dvchost=esmDestinationHost cn1Label=ESM Destination Port cn1=esmDestinationPort cs1Label=ConnectorName cs1=connectorName cs2Label=Connector Location cs2=connectorLocation cs3Label=Logger Location cs3=loggerLocation
Forwarders			
logger: 605	/Logger/Component/Forwarder/Configuration/Add	Forwarder [name] has been added	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger: 606	/Logger/Component/Forwarder/Configuration/Delete	Forwarder [name] has been deleted	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger: 607	/Logger/Component/Forwarder/Configuration/Update	Forwarder [name] has been updated	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:608	/Logger/Component/Forwarder/Configuration/Enable	Forwarder [name] has been enabled	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:609	/Logger/Component/Forwarder/Configuration/Disable	Forwarder [name] has been disabled	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:663	/Logger/Component/Forwarder/Configuration/Pause	Forwarder [name] has been paused	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter
logger:664	/Logger/Component/Forwarder/Configuration/Resume	Forwarder [name] has been resumed	fname=forwarderName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=forwarderType dvc=forwarderIpAddr dvchost=forwarderHostName cn1Label=Forwarder Port cn1=forwarderPort cs1Label=Forwarder Filter cs1=forwarderFilter

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
Receivers			
logger: 600	/Logger/Component/Receiver/Configuration/Add	Receiver [name] has been added	fname=receiverName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger: 601	/Logger/Component/Receiver/Configuration/Delete	Receiver [name] has been deleted	fname=receiverName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger: 602	/Logger/Component/Receiver/Configuration/Update	Receiver [name] has been updated	fname=receiverName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger: 603	/Logger/Component/Receiver/Configuration/Enable	Receiver [name] has been enabled	fname=receiverName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort
logger: 604	/Logger/Component/Receiver/Configuration/Disable	Receiver [name] has been disabled	fname=receiverName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=receiverType dvc=receiverIpAddr dvchost=receiverHostName cn1Label=Receiver Port cn1=receiverPort

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
SNMP Destinations			
logger: 644	/Logger/Component/SnmpDestination/Configuration/Add	SNMP destination [name] has been added	fname=snmpDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=SNMP Destination fileId=snmpDestinationId dvc=snmpDestinationIp dvchost=snmpDestinationHost cn1Label=SNMP Destination Port cn1=snmpDestinationPort cs1Label=ConnectorName cs1=connectorName cs2Label=Connector Location cs2=connectorLocation cs3Label=Logger Location cs3=loggerLocation
logger: 645	/Logger/Component/SnmpDestination/Configuration/Delete	SNMP destination [name] has been deleted	fname=snmpDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=SNMP Destination fileId=snmpDestinationId dvc=snmpDestinationIp dvchost=snmpDestinationHost cn1Label=SNMP Destination Port cn1=snmpDestinationPort cs1Label=ConnectorName cs1=connectorName cs2Label=Connector Location cs2=connectorLocation cs3Label=Logger Location cs3=loggerLocation
Syslog Destinations			
logger: 647	/Logger/Resource/SyslogDestination/Configuration/Add	Syslog destination [name] has been added	fname=syslogDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Syslog Destination fileId=syslogDestinationId dvc=syslogDestinationIp dvchost=syslogDestinationHost cn1Label=Syslog Destination Port cn1=syslogDestinationPort

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger: 648	/Logger/Component/SyslogDestination/Configuration/Delete	Syslog destination [name] has been deleted	fname=syslogDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Syslog Destination fileId=syslogDestinationId dvc=syslogDestinationIp dvchost=syslogDestinationHost cn1Label=Syslog Destination Port cn1=syslogDestinationPort
logger: 649	/Logger/Component/SyslogDestination/Configuration/Update	Syslog destination [name] has been updated	fname=syslogDestinationName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Syslog Destination fileId=syslogDestinationId dvc=syslogDestinationIp dvchost=syslogDestinationHost cn1Label=Syslog Destination Port cn1=syslogDestinationPort
Archives			
logger: 520	/Logger/Resource/Archive/Configuration/Add	Archive [archiveName] has been added	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 521	/Logger/Resource/Archive/Configuration/Delete	Archive [archiveName] has been deleted	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 523	/Logger/Resource/Archive/Configuration/Load	Archive [archiveName] has been loaded	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 524	/Logger/Resource/Archive/Configuration/Unload	Archive [archiveName] has been unloaded	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger: 525	/Logger/Resource/Archive/Configuration/Archive	Archive [archiveName] has been archived	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 526	/Logger/Resource/Archive/Add	Event archive settings added	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 527	/Logger/Resource/Archive/Update	Daily archive task settings updated	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
logger: 528	/Logger/Resource/Archive/Failed	Event archive failed	fname=archiveName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=EventArchive fileId=archiveId
Dashboards			
logger: 580	/Logger/Resource/Dashboard/Configuration/Add	Dashboard [name] has been added	fname=dashboardName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Dashboard fileId=DashboardId rt=receiptTime
logger: 581	/Logger/Resource/Dashboard/Configuration/Add	Dashboard [name] has been deleted	fname=dashboardName duser=UserName duid=userId cs4=sessionIdfile fileType=Dashboard fileId=DashboardId rt=receiptTime
logger: 582	/Logger/Resource/Dashboard/Configuration/Update	Dashboard [name] has been updated	fname=dashboardName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Dashboard fileId=DashboardId rt=receiptTime

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
Devices			
logger:510	/Logger/Resource/Device/Configuration/Add	Device [deviceName] has been added	fname=deviceName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Device fileId=deviceId
logger:511	/Logger/Resource/Device/Configuration/Delete	Device [deviceName] has been deleted	fname=deviceName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Device fileId=deviceId
logger:512	/Logger/Resource/Device/Configuration/Update	Device [deviceName] has been updated	fname=deviceName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Device fileId=deviceId
Filters			
logger:500	/Logger/Resource/Filter/Configuration/Add	Filter [filterName] has been added	fname=filterName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId
logger:501	/Logger/Resource/Filter/Configuration/Delete	Filter [filterName] has been deleted	fname=filterName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId
logger:502	/Logger/Resource/Filter/Configuration/Update	Filter [filterName] has been updated	fname=filterName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Filter fileId=filterId
Groups			
logger:513	/Logger/Resource/Group/Configuration/Add	Group [groupName] has been added	fname=groupName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Group fileId=groupId

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:514	/Logger/Resource/Group/Configuration/Delete	Group [groupName] has been deleted	fname=groupName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Group fileId=groupId
logger:515	/Logger/Resource/Group/Configuration/Update	Group [groupName] has been updated	fname=groupName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Group fileId=groupId
Peer Loggers			
logger:550	/Logger/Resource/PeerLogger/Configuration/Add	Peer Logger [name] has been added	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger fileId=LoggerId
logger:551	/Logger/Resource/PeerLogger/Configuration/Delete	Peer Logger [name] has been deleted	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger fileId=LoggerId
logger:570	/Logger/Resource/PeerLogger/Authorizations/Configuration/Add	Peer Logger authorization [name] has been added	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger Authorization
logger:571	/Logger/Resource/PeerLogger/Authorizations/Configuration/Delete	Peer Logger authorization [name] has been deleted	fname=Name duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Peer Logger Authorization fileId=LoggerId
Parsers			
logger:590	/Logger/Resource/ParserDescription/Configuration/Add	Parser Description [name] has been added	fileType=Parser Description duid=1 cs4=sessionIdfile cs4Label=Session ID duser=UserName rt=receiptTime fname=parserName

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger:591	/Logger/Resource/ParserDescription/Configuration/Delete	Parser Description [name] has been deleted	fileType=Parser Description cs4=sessionIdfile duser=UserName fileId=ParserID 710 duid=1 cs4Label=Session ID rt=receiptTime fname=parserName
logger:592	/Logger/Resource/ParserDescription/Configuration/Update	Parser Description [name] has been updated	fileType=Parser Description cs4=sessionIdfile duser=UserName fileId=ParserID duid=1 cs4Label=Session ID rt=receiptTime fname=parserName
Saved Searches			
logger:540	/Logger/Resource/SavedSearch/Configuration/Add	Saved search [name] has been added	fname=savedSearchName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
logger:541	/Logger/Resource/SavedSearch/Configuration/Delete	Saved search [name] has been deleted	fname=savedSearchName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
logger:542	/Logger/Resource/SavedSearch/Configuration/Update	Saved search [name] has been updated	fname=savedSearchName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Saved Search fileId=savedSearchId
Source Types			
logger:596	/Logger/Resource/SourceType/Configuration/Add	Source Type [name] has been added	cs4=sessionIdfile fileType=Source Type duid=1 cs4Label=Session ID duser=UserName rt=receiptTime fname=SourceTypeName
logger:597	/Logger/Resource/SourceType/Configuration/Delete	Source Type [name] has been deleted	fileType=Source Type cs4=sessionIdfile duser=UserName fileId=SourceTypeId duid=1 cs4Label=Session ID rt=receiptTime fname=SourceTypeName

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
logger: 598	/Logger/Resource/ /SourceType/Configuration/Update	Source Type [name] has been updated	fileType=Source Type cs4=sessionIdfile duser=UserName fileId=1SourceTypeId duid=1 cs4Label=Session ID rt=receiptTime fname=SourceTypeName
Storage Groups			
logger: 530	/Logger/Resource/ StorageGroup/Configuration/Add	Storage group [storageGroupName] has been added	fname=storageGroupName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Group fileId=storageGroupId
logger: 532	/Logger/Resource/ StorageGroup/Configuration/Update	Storage group [storageGroupName] has been updated	fname=storageGroupName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Group fileId=storageGroupId
Storage Rules			
logger: 533	/Logger/Resource/ StorageRule/Configuration/Add	Storage rule [name] has been added	fname=storageRuleName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Rule
logger: 535	/Logger/Resource/ StorageRule/Configuration/Update	Storage rule [name] has been updated	fname=storageRuleName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Rule
Storage Volume			
logger: 536	/Logger/Resource/ StorageVolume/Configuration/Add	Storage volume [name] has been added	fname=storageVolumeName duser=UserName duid=userId cs4=sessionIdfile cs4Label=Session ID fileType=Storage Volume fileId=storageVolumeId

Device Event Class ID	Device Event Category (cat)	Message	Additional Fields
Search			
logger: 680	/Logger/Search/Index/Update	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	/Logger/Search/Options/Update	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	/Logger/Search/Cancelled	Search session [sessionId] has been cancelled by [user]	cs1Label=Session ID duid=1 cs1=sessionIdfile duser=UserName rt=receiptTime
Maintenance Mode			
logger: 700	/Logger/Server/MaintenanceMode/Enter	Maintenance mode entered	fname=Maintenance Mode duser=UserName duid=userId cs4=sessionId cs4Label=Session ID fileType=Maintenance Mode

Appendix D

Examples of System Health Events

The following table provides examples of system health events generated on Logger. These examples are intended to help you understand the format and various fields of the generated events.



You can set up Alerts to be triggered to let you know when system health events are generated. [“Alerts” on page 296.](#)

The table includes information on the following system health event classes:

[“cpu” on page 581](#)
[“disk” on page 581](#)
[“eps” on page 582](#)
[“hardware” on page 582](#)
[“memory” on page 584](#)
[“network” on page 584](#)
[“raid” on page 584](#)
[“search” on page 585](#)
[“storagegroup” on page 585](#)

Device Event Class: ID	Example
cpu	
cpu: 100	CEF:0 ArcSight Logger 5.1.0.5780.0 cpu: 100 CPU Usage 1 cat=/Monitor/CPU/Usage cn1=3 cn1Label=Percent Usage cs2=CurrentValue cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302739080014 rt=1302739080014
disk	
disk: 101	CEF:0 ArcSight Logger 5.1.0.5803.0 disk: 101 Root Disk Space Remaining 1 cat=/Monitor/Disk/Space/Remaining/Root cn1=99 cn1Label=Percent Available cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=Ok cs4Label=Raw Status cs5=Root cs5Label=Location cs6=Disk/Space/Remaining/Root cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303927171790 rt=1303927171790

Device Event Class: ID	Example
disk: 102	CEF:0 ArcSight Logger 5.1.0.5780.0 disk: 102 Disk bytes read 1 cat=/Monitor/Disk/Read cn1=373524 cn1Label=Kb Read cs2=SinceStartup cs2Label=timeframe cs6=c0d0 cs6Label=Partition Name dst=192.168.35.115 dvc=192.168.35.115 end=1302743760036 rt=1302743760036
disk: 103	CEF:0 ArcSight Logger 5.1.0.5780.0 disk: 103 Disk bytes written 1 cat=/Monitor/Disk/Write cn1=24474998 cn1Label=Kb Written cs2=SinceStartup cs2Label=timeframe cs6=c0d0 cs6Label=Partition Name dst=192.168.35.115 dvc=192.168.35.115 end=1302743760038 rt=1302743760038
eps	
eps: 100	CEF:0 ArcSight Logger 5.1.0.5780.0 eps: 100 Overall Receiver EPS 1 cat=/Monitor/Receiver/All/EPS cn1=0 cn1Label=EPS cs2=SinceLastMonitorEvent cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302733680034 rt=1302733680034
eps: 101	CEF:0 ArcSight Logger 5.1.0.5780.0 eps: 101 Overall Forwarder EPS 1 cat=/Monitor/Forwarder/All/EPS cn1=0 cn1Label=EPS cs2=SinceLastMonitorEvent cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115
eps: 102	CEF:0 ArcSight Logger 5.1.0.5803.0 eps: 102 Individual Receiver EPS 1 cat=/Monitor/Receiver/EPS/Individual cn1=0 cn1Label=EPS cs1=N/A cs1Label=Receiver Type cs2=SinceLastMonitorEvent cs2Label=timeframe cs6=udp514 cs6Label=name dst=192.168.35.115 dvc=192.168.35.115 end=1303927500046 rt=1303927500046
eps: 103	CEF:0 ArcSight Logger 5.1.0.5780.0 eps: 103 Individual Forwarder EPS 1 cat=/Monitor/Forwarder/One/EPS cn1=0 cn1Label=EPS cs1=N/A cs1Label=Forwarder Type cs2=SinceLastMonitorEvent cs2Label=timeframe cs6=esm cs6Label=name dst=192.168.35.115 dvc=192.168.35.115 end=1302733080052 rt=1302733080052
hardware	
hardware: 101	CEF:0 ArcSight Logger 5.1.0.5784.0 hardware: 101 Electrical (Current) OK 1 cat=/Monitor/Sensor/Current/Ok cs1=0.80 Amps cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=10.2 (Power Supply) cs5Label=Location cs6=Current 2 cs6Label=Sensor Name dst=192.168.36.5 dvc=192.168.36.5 end=1303937520837 rt=1303937520837
hardware: 102	CEF:0 ArcSight Logger 5.1.0.5776.0 hardware: 102 Electrical (Current) Degraded 5 cat=/Monitor/Sensor/Current/Degraded cs1=126 Watts cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Inc cs4Label=Raw Status cs5Label=Location cs6=Power Meter cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302817019262 rt=1302817019262
hardware: 103	CEF:0 ArcSight Logger 5.1.0.5776.0 hardware: 103 Electrical (Current) Failed 8 cat=/Monitor/Sensor/Current/Failed cs1=126 Watts cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5Label=Location cs6=Power Meter cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302817019262 rt=1302817019262
hardware: 111	CEF:0 ArcSight Logger 5.1.0.5780.0 hardware: 111 Electrical (Voltage) OK 1 cat=/Monitor/Sensor/Voltage/Ok cs1=State Deasserted cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=3.1 (Processor) cs5Label=Location cs6=VCORE cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302819047959 rt=1302819047959

Device Event Class: ID	Example
hardware: 112	CEF:0 ArcSight Logger 5.1.0.5780.0 hardware: 112 Electrical (Voltage) Degraded 5 cat=/Monitor/Sensor/Voltage/Degraded cs1=State Deasserted cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Inc cs4Label=Raw Status cs5=3.1 (Processor) cs5Label=Location cs6=VCORE cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302819047959 rt=1302819047959
hardware: 113	CEF:0 ArcSight Logger 5.1.0.5780.0 hardware: 113 Electrical (Voltage) Failed 8 cat=/Monitor/Sensor/Voltage/Failed cs1=State Deasserted cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=Icr cs4Label=Raw Status cs5=3.1 (Processor) cs5Label=Location cs6=VCORE cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302819047959 rt=1302819047959
hardware: 121	CEF:0 ArcSight Logger 5.1.0.5803.0 hardware: 121 Battery OK 1 cat=/Monitor/Sensor/Battery/Ok cs1=cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=CMOS Battery cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303937972008 rt=1303937972008
hardware: 122	CEF:0 ArcSight Logger 5.1.0.5803.0 hardware: 122 Battery Degraded 5 cat=/Monitor/Sensor/Battery/Degraded cs1=cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Inc cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=CMOS Battery cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303937972008 rt=1303937972008
hardware: 123	CEF:0 ArcSight Logger 5.1.0.5803.0 hardware: 123 Battery Failed 8 cat=/Monitor/Sensor/Battery/Degraded cs1=cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=Icr cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=CMOS Battery cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303937972008 rt=1303937972008
hardware: 131	CEF:0 ArcSight Logger 5.1.0.5780.0 hardware: 131 Fan OK 1 cat=/Monitor/Sensor/Fan/Ok cs1=29.01 unspecified cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=Inc cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=Fan Block 1 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302823237825 rt=1302823237825
hardware: 132	
hardware: 133	CEF:0 ArcSight Logger 5.1.0.5780.0 hardware: 133 Fan Failure 8 cat=/Monitor/Sensor/Fan/Failed cs1=29.01 unspecified cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=Icr cs4Label=Raw Status cs5=7.1 (System Board) cs5Label=Location cs6=Fan Block 1 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302823237825 rt=1302823237825
hardware: 141	CEF:0 ArcSight Logger 5.1.0.5803.0 hardware: 141 Power Supply OK 1 cat=/Monitor/Sensor/PowerSupply/Ok cs1=cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=10.1 (Power Supply) cs5Label=Location cs6=Status cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1303938572149 rt=1303938572149
hardware: 142	

Device Event Class: ID	Example
hardware: 143	CEF:0 ArcSight Logger 5.1.0.5776.0 hardware: 143 Power Supply Failed 8 cat=/Monitor/Sensor/PowerSupply/Failed cs1=0 Watts cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=lcr cs4Label=Raw Status cs5=10.2 (Power Supply) cs5Label=Location cs6=Power Supply 2 cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302817019263 rt=1302817019263
hardware: 151	CEF:0 ArcSight Logger 5.1.0.5776.0 hardware: 151 Temperature OK 1 cat=/Monitor/Sensor/Temperature/Ok cs1=17 degrees C cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=ok cs4Label=Raw Status cs5=64.1 (Unknown (0x40)) cs5Label=Location cs6=Temp 1 cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302823560051 rt=1302823560051
hardware: 152	
hardware: 153	
memory	
memory: 100	CEF:0 ArcSight Logger 5.1.0.5780.0 memory: 100 Platform Memory Usage 1 cat=/Monitor/Memory/Usage/Platform cn1=2757 cn1Label=MB Used cs2=CurrentValue cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302797940018 rt=1302797940018
network	
network: 100	CEF:0 ArcSight Logger 5.1.0.5780.0 network: 100 Network Usage - Inbound 1 cat=/Monitor/Network/Usage/In cn1=41837428 cn1Label=Bytes Received cs2=SinceStartup cs2Label=timeframe cs6=eth0 cs6Label=Interface Name dst=192.168.35.115 dvc=192.168.35.115 end=1302733620026 rt=1302733620026
network: 101	CEF:0 ArcSight Logger 5.1.0.5780.0 network: 101 Network Usage - Outbound 1 cat=/Monitor/Network/Usage/Out cn1=158442791 cn1Label=Bytes Sent cs2=SinceStartup cs2Label=timeframe cs6=eth0 cs6Label=Interface Name dst=192.168.35.115 dvc=192.168.35.115 end=1302733620028 rt=1302733620028
raid	
raid: 101	CEF:0 ArcSight Logger 5.1.0.5780.0 raid: 101 RAID Controller OK 1 cat=/Monitor/RAID/Controller/Ok cs1=Type: RAID-5 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=Optimal cs4Label=Raw Status cs5=RAIDController/Configuration cs5Label=Location cs6=RAIDController/Configuration cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302886250104 rt=1302886250104
raid: 102	CEF:0 ArcSight Logger 5.1.0.5780.0 raid: 102 RAID Controller Degraded 5 cat=/Monitor/RAID/ControllerDegraded cs1=Type: RAID-5 Critical Disks: 0 Failed Disks: 0 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Degraded cs4Label=Raw Status cs5=RAIDController/Configuration cs5Label=Location cs6=RAIDController/Configuration cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302826128482 rt=1302826128482
raid: 103	

Device Event Class: ID	Example
raid: 111	CEF:0 ArcSight Logger 5.1.0.5776.0 raid: 111 RAID BBU OK 1 cat=/Monitor/RAID/BBU/Ok cs1=Battery/Capacitor Count: 1 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=OK cs4Label=Raw Status cs5=RAIDController/Battery/bbu cs5Label=Location cs6=RAIDController/Battery/bbu cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302890169285 rt=1302890169285
raid: 112	CEF:0 ArcSight Logger 5.1.0.5780.0 raid: 112 RAID BBU Degraded 5 cat=/Monitor/RAID/BBU/Degraded cs1=Fully Charged: false Remaining Time Alarm: false Remaining Capacity Alarm: false Over Charged: false isSOHGood: true cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Degraded cs3Label=Status cs4=Degraded cs4Label=Raw Status cs5=RAIDController/Battery/bbu cs5Label=Location cs6=RAIDController/Battery/bbu cs6Label=Sensor Name dst=192.168.35.115 dvc=192.168.35.115 end=1302820608015 rt=1302820608015
raid: 113	
raid: 121	CEF:0 ArcSight Logger 5.1.0.5780.0 raid: 121 RAID Disk OK 1 cat=/Monitor/RAID/DISK/Ok cs1=Port: 11 Box: 1 Bay: 1 Size: 500 GB Serial Number: 9SP24JD5 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Ok cs3Label=Status cs4=OK cs4Label=Raw Status cs5=Port: 11 Box: 1 Bay: 1 Serial Number: 9SP24JD5 cs5Label=Location cs6=RAIDController/Port/p0 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302849041777 rt=1302849041777
raid: 122	CEF:0 ArcSight Logger 5.1.0.5776.0 raid: 122 RAID Disk Rebuilding 5 cat=/Monitor/RAID/DISK/Rebuilding cs1=Port: 21 Box: 1 Bay: 1 Size: 1 TB Serial Number: WMATV6348517 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Rebuilding cs3Label=Status cs4=Rebuilding cs4Label=Raw Status cs5=Port: 21 Box: 1 Bay: 1 Serial Number: WMATV6348517 cs5Label=Location cs6=RAIDController/Port/p0 cs6Label=Sensor Name dst=192.168.36.37 dvc=192.168.36.37 end=1302826980530 rt=1302826980530
raid: 123	CEF:0 ArcSight Logger 5.1.0.5780.0 raid: 123 RAID Disk Failed 8 cat=/Monitor/RAID/DISK/Failed cs1=Port: 11 Box: 1 Bay: 2 Size: 500 GB Serial Number: 9SP23M08 cs1Label=Raw Value cs2=CurrentValue cs2Label=timeframe cs3=Failed cs3Label=Status cs4=Failed cs4Label=Raw Status cs5=Port: 11 Box: 1 Bay: 2 Serial Number: 9SP23M08 cs5Label=Location cs6=RAIDController/Port/p1 cs6Label=Sensor Name dst=192.168.37.21 dvc=192.168.37.21 end=1302826358346 rt=1302826358346
search	
search: 100	CEF:0 ArcSight Logger 5.1.0.5780.0 search: 100 Number of Searches Performed 1 cat=/Monitor/Search cn1=0 cn1Label=Number of Searches cs2=SinceStartup cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1302741300026 rt=1302741300026
storagegroup	
storagegroup: 100	CEF:0 ArcSight Logger 5.1.0.5803.0 storagegroup: 100 Storage Group Space Used 1 cat=/Monitor/StorageGroup/Space/Used cn1=1 cn1Label=Percent Used cn2=7 cn2Label=retention period (days) cn3=1024 cn3Label=used (MB) cs2=CurrentValue cs2Label=timeframe dst=192.168.35.115 dvc=192.168.35.115 end=1303928072008 fileType=storageGroup fname=Default Storage Group fsize=1534 rt=1303928072008

Appendix E

Event Field Name Mappings

The nomenclature for field names depends on the function area of the Logger. The following table provides the mapping between these types of names.

Database Name: Field name created in the database when you index this field. There will be no database name for a field if you have not indexed it. This field name is used when creating a SQL query for generating a report.

Search Results: Field name displayed in the search results when your search returns data in this field.

CEF Field Name: The key or field name as defined in Implementing ArcSight CEF. For a downloadable a copy of this guide, search for "ArcSight Common Event Format (CEF) Guide" on the Protect 724 Community at <https://protect724.arcsight.com>

Reports: Field name displayed in a report containing data from this field.

Database Name	Search Results	CEF Field Name	Reports
arc_agentAddress	agentAddress	agt	Agent Address
arc_agentHostName	agentHostName	ahost	Agent Host Name
arc_agentNtDomain	agentNtDomain	agentNtDomain	Agent NT Domain
arc_agentSeverity	agentSeverity	Severity	Severity
arc_agentType	agentType	at	Agent Type
arc_agentZone	agentZone	agentZone	Agent Zone
arc_agentZoneName	agentZoneName	agentZoneName	Agent Zone Name
arc_agentZoneResource	agentZoneResource	agentZoneResource	Agent Zone Resource
arc_agentZoneURI	agentZoneURI	agentZoneURI	Agent Zone URI
arc_applicationProtocol	applicationProtocol	app	Application Protocol
arc_baseEventCount	baseEventCount	cnt	Base Event Count
arc_bytesIn	bytesIn	in	Bytes In
arc_bytesOut	bytesOut	out	Bytes Out
arc_categoryBehavior	categoryBehavior	categoryBehavior	Category Behavior
arc_categoryDeviceGroup	categoryDeviceGroup	categoryDeviceGroup	Category Device Group

Database Name	Search Results	CEF Field Name	Reports
arc_categoryObject	categoryObject	categoryObject	Category Object
arc_categoryOutcome	categoryOutcome	categoryOutcome	Category Outcome
arc_categorySignificance	categorySignificance	categorySignificance	Category Significance
arc_categoryTechnique	categoryTechnique	categoryTechnique	Category Technique
arc_customerName	customerName	customerName	Customer Name
arc_destinationAddress	destinationAddress	dst	Destination Address
arc_destinationDnsDomain	destinationDnsDomain	destinationDnsDomain	Destination DNS Domain
arc_destinationHostName	destinationHostName	dhost	Destination Host Name
arc_destinationMacAddress	destinationMacAddress	dmac	Destination Mac Address
arc_destinationNtDomain	destinationNtDomain	dntdom	Destination NT Domain
arc_destinationPort	destinationPort	dpt	Destination Port
arc_destinationProcessName	destinationProcessName	dproc	Destination Process Name
arc_destinationServiceName	destinationServiceName	destinationServiceName	Destination Service Name
arc_destinationTranslatedAddress	destinationTranslatedAddress	destinationTranslatedAddress	Destination Translated Address
arc_destinationUserId	destinationUserId	duid	Destination User ID
arc_destinationUserName	destinationUserName	duser	Destination User Name
arc_destinationUserPrivileges	destinationUserPrivileges	dpriv	Destination User Privileges
arc_destinationZone	destinationZone	destinationZone	Destination Zone
arc_destinationZoneName	destinationZoneName	destinationZoneName	Destination Zone Name
arc_destinationZoneResource	destinationZoneResource	destinationZoneResource	Destination Zone Resource
arc_destinationZoneURI	destinationZoneURI	destinationZoneURI	Destination Zone URI
arc_deviceAction	deviceAction	act	Device Action
arc_deviceAddress	deviceAddress	dvc	Device Address
arc_deviceCustomDate1	deviceCustomDate1	deviceCustomDate1	Device Custom Date 1
arc_deviceCustomDate1Label	deviceCustomDate1Label	deviceCustomDate1Label	Device Custom Date 1 Label
arc_deviceCustomDate2	deviceCustomDate2	deviceCustomDate2	Device Custom Date 2
arc_deviceCustomDate2Label	deviceCustomDate2Label	deviceCustomDate2Label	Device Custom Date 2 Label

Database Name	Search Results	CEF Field Name	Reports
arc_deviceCustomNumber1	deviceCustomNumber1	cn1	Device Custom Number 1
arc_deviceCustomNumber1Label	deviceCustomNumber1Label	cn1Label	Device Custom Number 1 Label
arc_deviceCustomNumber2	deviceCustomNumber2	cn2	Device Custom Number 2
arc_deviceCustomNumber2Label	deviceCustomNumber2Label	cn2Label	Device Custom Number 2 Label
arc_deviceCustomNumber3	deviceCustomNumber3	cn3	Device Custom Number 3
arc_deviceCustomNumber3Label	deviceCustomNumber3Label	cn3Label	Device Custom Number 3 Label
arc_deviceCustomString1	deviceCustomString1	cs1	Device Custom String 1
arc_deviceCustomString1Label	deviceCustomString1Label	cs1Label	Device Custom String 1 Label
arc_deviceCustomString2	deviceCustomString2	cs2	Device Custom String 2
arc_deviceCustomString2Label	deviceCustomString2Label	cs2Label	Device Custom String 2 Label
arc_deviceCustomString3	deviceCustomString3	cs3	Device Custom String 3
arc_deviceCustomString3Label	deviceCustomString3Label	cs3Label	Device Custom String 3 Label
arc_deviceCustomString4	deviceCustomString4	cs4	Device Custom String 4
arc_deviceCustomString4Label	deviceCustomString4Label	cs4Label	Device Custom String 4 Label
arc_deviceCustomString5	deviceCustomString5	cs5	Device Custom String 5
arc_deviceCustomString5Label	deviceCustomString5Label	cs5Label	Device Custom String 5 Label
arc_deviceCustomString6	deviceCustomString6	cs6	Device Custom String 6
arc_deviceCustomString6Label	deviceCustomString6Label	cs6Label	Device Custom String 6 Label
arc_deviceEventCategory	deviceEventCategory	cat	Device Event Category
arc_deviceEventClassId	deviceEventClassId	Signature ID	Signature Id
arc_deviceExternalId	deviceExternalId	deviceExternalId	Device External Id
arc_deviceHostName	deviceHostName	dvchost	Device Host Name
arc_deviceInboundInterface	deviceInboundInterface	deviceInboundInterface	Device Inbound Interface

Database Name	Search Results	CEF Field Name	Reports
arc_deviceOutboundInterface	deviceOutboundInterface	deviceOutboundInterface	Device Outbound Interface
arc_deviceProduct	deviceProduct	Device Product	Device Product
arc_deviceReceiptTime	deviceReceiptTime	rt	Device Receipt Time
arc_deviceSeverity	deviceSeverity	deviceSeverity	Device Severity
arc_deviceVendor	deviceVendor	Device Vendor	Device Vendor
arc_deviceVersion	deviceVersion	Device Version	Device Version
arc_deviceZone	deviceZone	deviceZone	Device Zone
arc_deviceZoneName	deviceZoneName	deviceZoneName	Device Zone Name
arc_deviceZoneResource	deviceZoneResource	deviceZoneResource	Device Zone Resource
arc_deviceZoneURI	deviceZoneURI	deviceZoneURI	Device Zone URI
arc_endTime	endTime	end	End Time
arc_eventId	eventId	eventId	Event Id
arc_externalId	externalId	externalId	External Id
arc_fileName	fileName	fname	File Name
arc_filePath	filePath	filePath	File Path
arc_flexDate1	flexDate1	flexDate1	Flex Date 1
arc_flexDate1Label	flexDate1Label	flexDate1Label	Flex Date 1 Label
arc_flexNumber1	flexNumber1	flexNumber1	Flex Number1
arc_flexNumber1Label	flexNumber1Label	flexNumber1Label	Flex Number 1 Label
arc_flexNumber2	flexNumber2	flexNumber2	Flex Number 2
arc_flexNumber2Label	flexNumber2Label	flexNumber2Label	Flex Number 2 Label
arc_flexString1	flexString1	flexString1	Flex String 1
arc_flexString1Label	flexString1Label	flexString1Label	Flex String 1 Label
arc_flexString2	flexString2	flexString2	Flex String 2
arc_flexString2Label	flexString2Label	flexString2Label	Flex String 2 Label
arc_message	message	msg	Message
arc_name	name	Name	Name
arc_priority	priority	priority	Priority
arc_requestClientApplication	requestClientApplication	requestClientApplication	Request Client Application
arc_requestContext	requestContext	requestContext	Request Context
arc_requestMethod	requestMethod	requestMethod	Request Method
arc_requestUrl	requestUrl	request	Request URL
arc_requestUrlFilename	requestUrlFilename	requestUrlFilename	Request URL File Name

Database Name	Search Results	CEF Field Name	Reports
arc_requestUriQuery	requestUriQuery	requestUriQuery	Request URL Query
arc_sessionId	sessionId	sessionId	Session Id
arc_sourceAddress	sourceAddress	src	Source Address
arc_sourceHostName	sourceHostName	shost	Source Host Name
arc_sourceMacAddress	sourceMacAddress	smac	Source Mac Address
arc_sourceNtDomain	sourceNtDomain	sntdom	Source NT Domain
arc_sourcePort	sourcePort	spt	Source Port
arc_sourceProcessName	sourceProcessName	sproc	Source Process Name
arc_sourceServiceName	sourceServiceName	sourceServiceName	Source Service Name
arc_sourceTranslatedAddress	sourceTranslatedAddress	sourceTranslatedAddress	Source Translated Address
arc_sourceUserId	sourceUserId	suid	Source User Id
arc_sourceUserName	sourceUserName	suser	Source User Name
arc_sourceUserPrivileges	sourceUserPrivileges	spriv	Source User Privileges
arc_sourceZone	sourceZone	sourceZone	Source Zone
arc_sourceZoneName	sourceZoneName	sourceZoneName	Source Zone Name
arc_sourceZoneResource	sourcezoneResource	sourceZoneResource	Source Zone Resource
arc_sourceZoneURI	sourceZoneURI	sourceZoneURI	Source Zone URI
arc_startTime	startTime	start	Start Time
arc_transportProtocol	transportProtocol	proto	Transport Protocol
arc_type	type	type	Type
arc_vulnerabilityExternalID	vulnerabilityExternalID	vulnerabilityExternalID	Vulnerability External Id
arc_vulnerabilityURI	VulnerabilityURI	vulnerabilityURI	Vulnerability URI

Appendix F

Logger Content

This chapter contains information about the following Logger content:

- “Reports” on page 593
- “Parameters” on page 626
- “System Filters” on page 631

Reports

Logger provides the reports described in the tables below. In the Logger UI, these reports are listed in categories, accessible through the Category Explorer (in the left pane). For example, the [Errors Detected in Anti-Virus Deployment](#) report is listed in the [Anti-Virus](#) category and the [Anti-Virus](#) category is listed in the parent category called [Device Monitoring](#).

The reports contain hyperlinks that drill-down to other reports. For example, the [Most Common Events by Severity](#) report displays a field called `Count`. Clicking on the `Count` field drills down to the [Destination Counts by Device Severity](#) report, which provides additional detail information, as shown in the following figure. The drill-down relationship between reports is shown in the tables below.

Severity	Name	Count
Very-High	SMB: WinLogon DoS	2089
Very-High	IRC: Trojan.IrcBounce Command Channel	10
Very-High	Snort Alarm [1:2404:5]	58
Very-High	Host is DOWN	41
Very-High	ids syn attack	26
Very-High	RCRS/POP3_INVALID_ARG_TO_QUIT	16

Severity	Target Zone	Target Address	Count
Very-High	/All Zones/ArcSight System/Private Address Space Zones/RFC1918: 10.0.0.0-10.255.255.255	10.0.0.10	57
Very-High	/All Zones/ArcSight System/Private Address Space Zones/RFC1918: 10.0.0.0-10.255.255.255	10.156.106.101	25
Very-High	/All Zones/ArcNet Zones/sj2.west.arconet.com - internal	10.0.112.187	10
Very-High	/All Zones/ArcSight System/Private Address Space Zones/RFC1918: 10.0.0.0-10.255.255.255	10.0.0.10	10
Very-High	/All Zones/ArcSight System/Private Address Space Zones/RFC1918: 172.16.0.0-172.31.255.255	172.16.5.157	8

Logger provides the following top-level categories:

- “Device Monitoring” on page 594
- “Foundation” on page 604

- [“SANS Top 5” on page 619](#)

Device Monitoring

This category provides a device or application based view on events.

The following categories are located under the Device Monitoring category:

- [“Anti-Virus” on page 594](#)
- [“CrossDevice” on page 595](#)
- [“Database” on page 599](#)
- [“Firewall” on page 599](#)
- [“IDS-IPS” on page 600](#)
- [“Identity Management” on page 601](#)
- [“Network” on page 601](#)
- [“Operating System” on page 602](#)
- [“VPN” on page 603](#)

Anti-Virus

This is a sub-category of the Device Monitoring category, focusing on events related to Anti-Virus systems.

The Anti-Virus category is located under the following path.

Device Monitoring\Anti-Virus

The Anti-Virus category reports are listed in the following table.

Table F-1 Anti-Virus Category Reports

Report	Description	Drill Down	Parameters
Errors Detected in Anti-Virus Deployment	This report shows a summary of information on the anti-virus errors, including the Anti-Virus product information, host details, error information, and the number of errors.	none	none
Failed Anti-Virus Updates	This report displays a table with the Device Vendor, Device Product Target Zone Name, Target Host Name, Target Address, and Minute(EndTime).	none	none
Top Infected Systems	This report displays summaries of the systems reporting the most infections.	none	none
Update Summary	This report shows a summary and details of anti-virus update activity.	none	none

Report	Description	Drill Down	Parameters
Virus Activity by Hour	This report shows malware activity by hour.	none	none

CrossDevice

This is a sub-category of the Device Monitoring category. It provides information on events that are similar across devices, e.g., logins, start up and shut down, etc.

The CrossDevice category is located under the following path.

Device Monitoring\CrossDevice

The CrossDevice category reports are listed in the following table.

Table F-2 CrossDevice Category Reports

Report	Description	Drill Down	Parameters
Bandwidth Usage by Hour	This report shows the network bandwidth usage per hour by device type. Note that the bandwidth values are based on all reported traffic, including traffic within the network, as well as traffic to and from internal and external sources. There is a parameter allowing the limitation of the devices to one of the following: Firewalls, Network Equipment, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: "Bandwidth Usage by Hour" on page 595 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "dmBandwidthParameter" on page 629
Bandwidth Usage by Protocol	This report shows all the protocols sorted by bandwidth usage, by device type. Note that the bandwidth values are based on all reported traffic, including traffic within the network, as well as traffic to and from internal and external sources. There is a parameter allowing the limitation of the devices to one of the following: Firewalls, Network Equipment, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: "Bandwidth Usage by Protocol" on page 595 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "dmBandwidthParameter" on page 629

Report	Description	Drill Down	Parameters
By User Account - Accounts Created	This report shows all newly created accounts that were reported to Logger.	none	none
Configuration Changes by Type	This report shows recent configuration changes that were reported to Logger.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: "Configuration Changes by Type" on page 596 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "dmConfigurationParameter" on page 629
Configuration Changes by User	This report shows recent configuration changes that were reported to Logger.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: "Configuration Changes by User" on page 596 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "dmConfigurationParameter" on page 629
Failed Login Attempts	This report shows authentication failures from login attempts by hour. There is a parameter allowing the limitation of the devices to one of the following: Database, Firewalls, Identity Management systems, Network Equipment, Operating Systems, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: "Failed Login Attempts" on page 596 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "dmLoginParameter" on page 629
Failed Logins by Destination Address	This report shows authentication failures from login attempts by destination address. There is a parameter allowing the limitation of the devices to one of the following: Database, Firewalls, Identity Management systems, Network Equipment, Operating Systems, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: "Failed Logins by Destination Address" on page 596 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "dmLoginParameter" on page 629

Report	Description	Drill Down	Parameters
Failed Logins by Source Address	This report shows authentication failures from login attempts by source address. There is a parameter allowing the limitation of the devices to one of the following: Database, Firewalls, Identity Management systems, Network Equipment, Operating Systems, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: “Failed Logins by Source Address” on page 597 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “dmLoginParameter” on page 629
Failed Logins by User	This report shows authentication failures from login attempts by user. There is a parameter allowing the limitation of the devices to one of the following: Database, Firewalls, Identity Management systems, Network Equipment, Operating Systems, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: “Failed Logins by User” on page 597 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “dmLoginParameter” on page 629
Login Event Audit	This report shows all authentication events. There is a parameter allowing the limitation of the devices to one of the following: Database, Firewalls, Identity Management systems, Network Equipment, Operating Systems, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: “Login Event Audit” on page 597 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “dmLoginParameter” on page 629
Password Changes	This report shows all password changes that were reported to Logger.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: “Password Changes” on page 597 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “dmLoginParameter” on page 629

Report	Description	Drill Down	Parameters
Successful Logins by Destination Address	This report shows successful authentication events by destination addresses. There is a parameter allowing the limitation of the devices to one of the following: Database, Firewalls, Identity Management systems, Network Equipment, Operating Systems, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: “Successful Logins by Destination Address” on page 598 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “dmLoginParameter” on page 629
Successful Logins by Source Address	This report shows successful authentication events by source addresses. There is a parameter allowing the limitation of the devices to one of the following: Database, Firewalls, Identity Management systems, Network Equipment, Operating Systems, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: “Successful Logins by Source Address” on page 598 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “dmLoginParameter” on page 629
Successful Logins by User	This report shows successful authentication events by user. There is a parameter allowing the limitation of the devices to one of the following: Database, Firewalls, Identity Management systems, Network Equipment, Operating Systems, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: “Successful Logins by User” on page 598 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “dmLoginParameter” on page 629

Report	Description	Drill Down	Parameters
Top Bandwidth Hosts	This report shows the top hosts, sorted by bandwidth usage. Note that the bandwidth values are based on all reported traffic, including traffic within the network, as well as traffic to and from internal and external sources. There is a parameter allowing the limitation of the devices to one of the following: Firewalls, Network Equipment, or VPNs. By default, the parameter is set to all devices and applications that report data.	none	This report prompts you to supply a value for the following parameter. <ul style="list-style-type: none"> “dmBandwidthParameter” on page 629
Top Hosts by Number of Connections	This report shows a summary of the number of connections by the top hosts. There is a parameter allowing the limitation of the devices to one of the following: Firewalls, Network Equipment, or VPNs. By default, the parameter is set to all devices and applications that report data.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Reporting Device field drill downs to the following report: “Top Hosts by Number of Connections” on page 599 <p>This report drills down to itself.</p>	This report prompts you to supply a value for the following parameter. <ul style="list-style-type: none"> “dmBandwidthParameter” on page 629

Database

This is a sub-category of the Cross Device category, focusing on database events.

The Database category is located under the following path.

Device Monitoring\Database

The Database category reports are listed in the following table.

Table F-3 Database Category Reports

Report	Description	Drill Down	Parameters
Database Errors and Warnings	This report shows recent database errors and warnings.	none	none

Firewall

This is a sub-category of the Device Monitoring category, focusing on firewall events.

The Firewall category is located under the following path.

Device Monitoring\Firewall

The Firewall category reports are listed in the following table.

Table F-4 Firewall Category Reports

Report	Description	Drill Down	Parameters
Denied Connections by Address	This report shows a summary and details of inbound and outbound connections denied by Firewall devices.	none	none
Denied Connections by Port	This report shows a summary and details of inbound and outbound ports denied by Firewall devices.	none	none
Denied Connections per Hour	This report shows a summary and details of inbound and outbound connections denied by Firewall devices on an hourly basis.	none	none

IDS-IPS

This is a sub-category of the Device Monitoring category, focusing on Intrusion Detection System and Intrusion Prevention System events.

The IDS-IPS category is located under the following path.

Device Monitoring\IDS-IPS

The IDS-IPS category reports are listed in the following table.

Table F-5 IDS-IPS Category Reports

Report	Description	Drill Down	Parameters
Alert Counts by Device	This report shows counts of IDS and IPS alerts.	none	none
Alert Counts by Port	This report shows count of IDS and IPS alerts by destination port.	none	none
Alert Counts by Severity	This report shows count of IDS and IPS alerts by agent severity.	none	none
Alert Counts by Type	This report shows the count of IDS and IPS alerts by type (category technique).	none	none
Alert Counts per Hour	This report shows the count of IDS and IPS alerts for each hour.	none	none
Top Alert Destinations	This report shows the top destinations of IDS and IPS alerts.	none	none

Report	Description	Drill Down	Parameters
Top Alerts from IDS and IPS	This report shows the top alerts coming from Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS).	none	none
Top Alert Sources	This report shows the top sources of IDS and IPS alerts.	none	none
Worm Infected Systems	This report shows a list of systems that have been infected by a worm.	none	none

Identity Management

This is a sub-category of the Device Monitoring category, focusing on Identity Management system events.

The Identity Management category is located under the following path.

`Device Monitoring\Identity Management`

The Identity Management category reports are listed in the following table.

Table F-6 Identity Management Category Reports

Report	Description	Drill Down	Parameters
Connection Counts by User	This reports shows count information about connections for each user reported by Identity Management devices.	none	none

Network

This is a sub-category of the Device Monitoring category, focusing on network devices such as routers and switches.

The Network category is located under the following path.

`Device Monitoring\Network`

The Network category reports are listed in the following table.

Table F-7 Network Category Reports

Report	Description	Drill Down	Parameters
Device Critical Events	This report shows information regarding critical events on network devices. These critical events could be indications of hardware failure, resource exhaustion, configuration issues, or attacks.	none	none
Device Errors	This report shows information regarding error events on network devices. These critical events could be indications of hardware failure, resource exhaustion, configuration issues, or attacks.	none	none
Device Events	This report shows information regarding events on network devices. These critical events could be indications of hardware failure, resource exhaustion, configuration issues, or attacks.	none	none
Device Interface Down Notifications	This report shows a table displaying the network devices that report a down link.	none	none
Device Interface Status Messages	This report shows the network devices reporting link status changes.	none	none
Device SNMP Authentication Failures	This report shows information regarding network device SNMP failures.	none	none

Operating System

This is a sub-category of the Device Monitoring category, focusing on operating system events.

The Operating System category is located under the following path.

Device Monitoring\Operating System

The Operating System category reports are listed in the following table.

Table F-8 Operating System Category Reports

Report	Description	Drill Down	Parameters
Login Errors by User	This report shows the details of failed logins for each username (time, event name, source, and destination).	none	none
User Administration	This report shows user and user group creations, modifications, and deletions.	none	none

VPN

This is a sub-category of the Device Monitoring category, focusing on virtual private network events.

The VPN category is located under the following path.

Device Monitoring\VPN

The VPN category reports are listed in the following table.

Table F-9 VPN Category Reports

Report	Description	Drill Down	Parameters
Authentication Errors	This report shows errors generated by a VPN connection attempt. The address is the IP address of the VPN connection source. This report can be used to see which users are having difficulties using or setting up their VPN clients.	none	none
Connection Counts by User	This report shows count information about VPN connections for each user. Details of each user's connection counts are provided, including connection count and systems accessed.	none	none
Connections Accepted by Address	This report shows successful VPN connection data.	none	none
Connections Denied by Address	This report shows denied VPN connection data.	none	none
Connections Denied by Hour	This report shows denied VPN connection data for each hour.	none	none

Foundation

This category covers a broad range of events, from security and perimeter defense to network bandwidth usage and configuration events.

The following categories are located under the Foundation category:

- [“Attackers” on page 610](#)
- [“Resource Access” on page 612](#)
- [“Targets” on page 614](#)
- [“User Tracking” on page 616](#)
- [“NetFlow Monitoring” on page 616](#)
- [“Network Monitoring” on page 617](#)
- [“Network Monitoring” on page 617](#)

Configuration Monitoring

This category covers configuration changes to systems and applications.

The Configuration Monitoring category is located under the following path.

Foundation\Configuration Monitoring

The Configuration Monitoring category reports are listed in the following table.

Table F-10 Configuration Monitoring Category Reports

Report	Description	Drill Down	Parameters
Accounts Created by User Account	This report details the successfully created accounts created on network hosts. The table includes the timestamp of when the account was created, the created account name (Destination User Name), the name of the user creating the account (Source User Name), the account creation event name, and the zone and host name of the device on which the account was created.	none	none
Accounts Deleted by Host	This report provides a listing of user deletions, ordered by Customer, Zone, and System.	none	none

Report	Description	Drill Down	Parameters
Accounts Deleted by User Account	This report displays a table showing the date, the deleted user name, the user name that deleted the account, the account deletion event name, and the zone and host name of the system from which the account was deleted.	none	none
Anti-Virus Updates-All-Failed	This report displays a table with the Device Vendor, Device Product Target Zone Name, Target Host Name, Target Address and Minute(EndTime), of all failed anti-virus update events.	none	none
Anti-Virus Updates-All-Summary	This report displays a table showing the Target Zone Name, Target Host Name, Target Address, Device Vendor, Device Product, Category Outcome and the sum of the Aggregated Event Count of all anti-virus events.	none	none
Asset Startup and Shutdown Event Log	This report provides a listing of the system startup and shutdown events.	none	none
Device Configuration Changes	This report shows a table of events related to successful device configuration modification events. The information provided includes the Device Group, the Zone, Address and Host Name, the Configuration Event name, the User ID and Name making the change, and the hour the change was made. Clicking on a device entry in the Device Group column will bring up a new report, Device Configuration Changes Drilldown, which will show only configuration events for that particular device type.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Device Group field drill downs to the following report: "Device Configuration Changes" on page 605 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "deviceGroupParameter" on page 628

Report	Description	Drill Down	Parameters
Device Configuration Events	This report shows a table of events related to various device configuration modification events, whether successful or not. The information provided includes the Device Group, the Zone, Address and Host Name, the Configuration Event name, the User ID and Name making the change, and the hour the change was made. Clicking on a device entry in the Device Group column will bring up a new report, Device Configuration Events Drilldown, which will show only configuration events for that particular device type.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Device Group field drill downs to the following report: “Device Configuration Events” on page 606 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “deviceGroupParameter” on page 628
Device Misconfigurations	This report shows a table of events related to device configuration checks. The information provided includes the Device Group, the Zone, Address and Host Name, the Misconfiguration name, and the count of the number of misconfigurations found. Clicking on a Device Group entry will run the Device Misconfigurations Drilldown report, focusing on the device type that was clicked.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Device Group field drill downs to the following report: “Device Misconfigurations” on page 606 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “deviceGroupParameter” on page 628
Password Changes	This report displays a table of user accounts having their passwords changed. The table shows the time the password was changed, the user name of the account with the new password, the zone and address of the system on which the password was changed, and the zone and address from which the change originated.	none	none
Vulnerability Scanner Logs by Host	This report shows Vulnerability Scanner Logs grouped by Zone and Host IP Address.	none	<p>This report prompts you to supply values for the following parameters.</p> <ul style="list-style-type: none"> “deviceProduct” on page 628 “deviceVendor” on page 629

Report	Description	Drill Down	Parameters
Vulnerability Scanner Logs by Vulnerability	This report shows Vulnerability Scanner Logs grouped by Vulnerability IDs and Names.	none	<p>This report prompts you to supply values for the following parameters.</p> <ul style="list-style-type: none"> • “deviceProduct” on page 628 • “deviceVendor” on page 629

Intrusion Monitoring

This is a sub-category of the Foundation category, focusing on security, perimeter defense, resource access and user tracking events.

The Intrusion Monitoring category is located under the following path.

Foundation\Intrusion Monitoring

The Intrusion Monitoring category reports are listed in the following table.

Table F-11 Intrusion Monitoring Category Reports

Report	Description	Drill Down	Parameters
Device Interface Down Notifications	This report shows a table displaying the network devices that report a down link.	none	none
Firewall Traffic by Service	This report displays a table showing the Port, transport protocol, application protocol, and the number of events reported by firewalls.	none	none
Least Common Events	This report displays all events in the time period selected and orders them by the sum of the aggregated event count in ascending order. The columns are hyperlinked for convenience. The Event Name column will bring up the Bottom Destinations report using the same time frame. The Count column will bring up the Bottom Sources report using the same time frame.	<p>This report provides drilldowns from the following fields.</p> <ul style="list-style-type: none"> • Event Name field drill downs to the following report: “Bottom Destinations” on page 614 • Count field drill downs to the following report: “Bottom Sources” on page 610 	none

Report	Description	Drill Down	Parameters
Most Common Events	This report displays the 200 most common events within the time range specified. The event name is hyperlinked to drilldown to the Destination Counts by Event Name report, which will show destination information for the event selected. The Count field will bring up the Source Counts by Destination Port report, which will include information about all sources by destination port.	This report provides drilldowns from the following fields. <ul style="list-style-type: none"> arc_name field drill downs to the following report: "Destination Counts by Event Name" on page 615 SUM(events.arc_baseEventCount) field drill downs to the following report: "Source Counts by Destination Port" on page 610 	none
Most Common Events by Severity	This report displays a table showing the Severity, event name and count of events in descending order.	This report provides drilldowns from the following fields. <ul style="list-style-type: none"> Severity field drill downs to the following report: "Source Counts by Device Severity" on page 611 Count field drill downs to the following report: "Destination Counts by Device Severity" on page 615 	none
Probes on Blocked Ports by Source	This report displays a table of events showing the source zone, address and host name, the transport protocol, the destination port, and the count of events where the destination port is in the list of commonly blocked ports. The query uses the comonlyblockedPorts parameter, which can be edited to include other ports (please make a copy of the report, the query, and the parameter, and modify your version as updates to the Foundation Content may overwrite your changes).	none	This report prompts you to supply a value for the following parameter. <ul style="list-style-type: none"> "commonlyBlockedPorts" on page 627
SecurityDashboardRpt	This custom report displays a table showing the source address, category behavior, destination address and event ID.	none	none

Report	Description	Drill Down	Parameters
SecurityDBR report	This custom Security Dashboard report displays two charts and a table. The first chart shows the number of events by source address. The second chart shows the number of events by destination address. The table shows the counts of events for each source and destination.	none	none
Top IDS Attack Events	This report displays a table showing the top events from IDS systems, with the IDS Event name, the type of IDS, and the count of each IDS event where the category significance is Compromise or Hostile.	none	none
Top IDS Events	This report displays a table showing the top events from IDS systems, with the IDS Event name, the type of IDS, and the count of each IDS event.	none	none
Top Machines Traversing Firewall	This report displays the source zone, address and hostname, and number of events reported by firewalls.	none	none
Top Web Traffic	This report displays a table showing the hour, source zone, address and host name, the web port and the count of events where the destination port is listed in the webPorts parameter.	none	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> • "webPorts" on page 630
Windows Events	This report displays a table showing the device zone, address and host name, the device event ID, the source user ID, user name and NT domain, the destination user ID, user name and NT domain, the behavior, outcome and event type, and the count of events of each type reported by any Microsoft operating system.	none	none
Worm Infected Systems	This report displays a table showing the Zone Name, Host Name and Address of systems exhibiting symptoms of being infected by a worm.	none	none

Attackers

This is a sub-category of the Intrusion Monitoring category, focusing on events based on source or attacker information.

The Attackers category is located under the following path.

Foundation\Intrusion Monitoring\Attackers

The Attackers category reports are listed in the following table.

Table F-12 Attackers Category Reports

Report	Description	Drill Down	Parameters
Bottom Sources	This report displays the Source Zone Names, Source Addresses and event Count ordered by the sum of the base event counts in ascending order. Clicking on the hyperlink for the Count column will bring up the Bottom Targets report. It is the target of the Least Common Events report's Count column.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Count field drill downs to the following report: "Bottom Targets" on page 615 	none
Source Counts by Destination	This report displays a table showing the destination zone and address, the source zone and the number of each event for a specific destination zone and address where the category significance is Compromise or Hostile.	<p>This report provides drilldowns from the following fields.</p> <ul style="list-style-type: none"> Destination Zone field drill downs to the following report: "Source Counts by Destination" on page 610 Destination Address field drill downs to the following report: "Source Counts by Destination" on page 610 Source Count field drill downs to the following report: "Attack Events by Destination" on page 614 <p>This report drills down to itself.</p>	<p>This report prompts you to supply values for the following parameters.</p> <ul style="list-style-type: none"> "destinationAddress" on page 628 "zones" on page 631
Source Counts by Destination Port	This report displays a table showing the Destination Port, the source zone and address, and the number of events for each port.	none	none

Report	Description	Drill Down	Parameters
Source Counts by Device	This report displays a table showing the device zone and address, the source zone and address, and the number of each event where the category significance is Compromise or Hostile.	none	none
Source Counts by Device Severity	This report displays a table showing the Severity, source zone and address, and the number of events at that severity.	none	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> • “deviceSeverityParameter” on page 628
Source Counts by Source Port	This report displays a table showing the Source Port, source zone and address, and a count of events where the category significance is Compromise or Hostile.	none	none
Source Port Counts	This report displays a table showing the Source Port, Event Name and count of the events where the category significance is Compromise or Hostile.	none	none
Top 10 Talkers	This report displays a table of the Top 10 systems generating events, showing the Source zone and address, as well as the number of events from that system.	none	none
Top Attacker Detail	This report displays a table showing the severity, attacker zone and address, the target zone and address, and the count of events for a specified source zone and address where the category significance is Compromise or Hostile.	none	none
Top Attacker Details	This report displays the Severity, Attacker Zone, Attacker Address, Target Zone, Target Address and Count of events where the Category Significance of the event is compromise or hostile, ordered by the event count in descending order. This report is the target of the Top Attackers report's Attacker Address column.	none	<p>This report prompts you to supply values for the following parameters.</p> <ul style="list-style-type: none"> • “IPAddress” on page 627 • “zoneParameter” on page 630

Report	Description	Drill Down	Parameters
Top Attacker Ports	This report displays a table showing the Attacker Port, Transport Protocol and the count of events where the category significance is Compromise or Hostile.	none	none
Top Attackers	This report shows the Attacker Zone Names, Attacker Addresses and Count of events where the Category Significance of the events is compromise or hostile, in descending order of the sum of the base event count. This report has hyperlinks that will run reports showing more information base on the field selected. The Attacker Zone column will run the Top Attack Sources report. The Attacker Address will run the Top Attacker Details report. The Count column will run the Top Targets report.	<p>This report provides drilldowns from the following fields.</p> <ul style="list-style-type: none"> Attacker Address field drill downs to the following report: "Top Attacker Details" on page 611 Count field drill downs to the following report: "Top Targets" on page 616 	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "zoneParameter" on page 630
Top Attack Sources	This report displays the Attacker Zone and Count of events where the Category Significance of the event is compromise or hostile, ordered by the event count in descending order. This report has a hyperlink in the Attacker Zone column that will run the Top Attackers report.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Attacker Zone field drill downs to the following report: "Top Attackers" on page 612 	none
Top Sources Detected by Snort	This report displays a table showing the source zone, address and host name and the number of events detected by Snort.	none	none
Top Sources Traversing Firewalls	This report displays a table of the source zone, address and host name, as well as the count of events, where the event was reported by a firewall device.	none	none

Resource Access

This is a sub-category of the Intrusion Monitoring category, focusing on protected resources.

The Resource Access category is located under the following path.

Foundation\Intrusion Monitoring\Resource Access

The Resource Access category reports are listed in the following table.

Table F-13 Resource Access Category Reports

Report	Description	Drill Down	Parameters
Access Events by Resource	This report displays a table showing the Resource Type, the zone and address, the access event, the outcome and the number of times this has happened over the time period selected. Clicking on a resource type will run the Access Events by Resource Drilldown report showing the events for the selected resource type.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Resource Type field drill downs to the following report: "Access Events by Resource" on page 613 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "categoryObjectParameter" on page 627
Least Common Accessed Ports	This report displays a table showing the Destination Port, the Transport Protocol and a count of the events for that port where the transport protocol is TCP or UDP.	none	none
Resource Access by Users - Failures	This report displays a table showing the Resource Type, the User ID and Name, Destination zone and address, the Access Event and number of failed attempts to access a given resource. Clicking on an entry in the Resource Type column will bring up the Resource Access by Users - Failures Drilldown report, which will show all related events for that resource type.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Resource Type field drill downs to the following report: "Resource Access by Users - Failures" on page 613 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "categoryObjectParameter" on page 627
Resource Access by Users - Successes-Attempts	This report displays a table showing the Resource Type, Outcome, destination user ID and name, destination zone and address, the access event name and the number of such events. The Resource Type column is hyperlinked so that clicking on a resource type will run the Resource Access by Users - Successes-Attempts Drilldown report, showing only events for the selected resource type.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Resource Type field drill downs to the following report: "Resource Access by Users - Successes-Attempts" on page 613 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> "categoryObjectParameter" on page 627

Report	Description	Drill Down	Parameters
Top Machines Accessing the Web	This report displays a table showing the source zone, address and host name, the destination port and the number of events where the destination port is in the webPorts parameter list.	none	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> • “webPorts” on page 630

Targets

This is a sub-category of the Intrusion Monitoring category, focusing on events based on destination or target information.

The Targets category is located under the following path.

Foundation\Intrusion Monitoring\Targets

The Targets category reports are listed in the following table.

Table F-14 Targets Category Reports

Report	Description	Drill Down	Parameters
Attack Events by Destination	This report displays a table showing the destination zone and address, the source zone and address, the event name and the number of each event for a specific destination zone and address where the category significance is Compromise or Hostile.	<p>This report provides drilldowns from the following fields.</p> <ul style="list-style-type: none"> • Destination Zone field drill downs to the following report: “Attack Events by Destination” on page 614 • Destination Address field drill downs to the following report: “Attack Events by Destination” on page 614 <p>This report drills down to itself.</p>	<p>This report prompts you to supply values for the following parameters.</p> <ul style="list-style-type: none"> • “destinationAddress” on page 628 • “zones” on page 631
Bottom Destinations	This report displays the Destination Zone Names, Destination Addresses and event Count ordered by the sum of the base event counts in ascending order. It is the target of the Least Common Events report's Event Name column.	none	none

Report	Description	Drill Down	Parameters
Bottom Targets	This report shows the Target Zone Names, Target Addresses and Count of events where the Category Significance of the events is compromise or hostile, in ascending order of the sum of the base event count. This report is the target of the Bottom Sources report's Count column.	none	none
Destination Counts by Device Severity	This report displays a table showing the Severity, target zone and address, and the number of events for each severity.	none	This report prompts you to supply a value for the following parameter. <ul style="list-style-type: none"> • "deviceSeverityParameter" on page 628
Destination Counts by Event Name	This report displays a table showing the event name, the target zone and address, and the number of events for each destination.	none	This report prompts you to supply a value for the following parameter. <ul style="list-style-type: none"> • "eventNameParameter" on page 630
Target Attack Counts by Severity	This report displays a table showing the Severity, the target zone and the number of events where the event has a category significance of Compromise, Hostile or Suspicious.	none	none
Target Counts by Event Name	This report displays a table showing the event name, target zone and address, and the number of time that event has occurred where the category significance is Compromise or Hostile.	none	none
Target Counts by Severity	This report displays a table showing the Severity, the target zone and address, and the number of events where the event has a category significance of Compromise, Hostile or Suspicious.	none	none
Target Counts by Source	This report displays a table showing the Source zone and address, the target zone and address, and the number of events where the event has a category significance of Compromise, Hostile or Suspicious.	none	none

Report	Description	Drill Down	Parameters
Target Counts by Source Port	This report displays a table showing the Source Port, the count of events for that port, with the destination zone and address, where the category significance is Compromise or Hostile.	none	none
Target Counts by Target Port	This report displays a table showing the Destination Port, the number of events for each port, and the target zone and address for events with category significance of Compromise or Hostile.	none	none
Target Port Counts	This report displays a table showing the Target Port, the number of events for that port, and the target zone and address of events where the category significance is Compromise or Hostile.	none	none
Top Destination Ports	This report displays a table of the top destination ports and the number of events for each port.	none	none
Top Destinations Across Firewalls	This report displays a table of the destination zone, address and host name, as well as the count of events, where the event was reported by a firewall device.	none	none
Top Destinations in IDS Events	This report displays a table showing the Destination zone, address and host name, as well as the count of event going to each host, for all events coming from an IDS.	none	none
Top Targets	This report displays the Target Zone, Target Address and Count of events where the Category Significance of the event is compromise or hostile, ordered by the event count in descending order. This report is the target of the Top Attackers report's Count column.	none	<p>This report prompts you to supply values for the following parameters.</p> <ul style="list-style-type: none"> • "IPAddress" on page 627 • "zoneParameter" on page 630

User Tracking

NetFlow Monitoring

This is a sub-category of the Foundation category, focusing on NetFlow data.

The NetFlow Monitoring category is located under the following path.

Foundation\NetFlow Monitoring

The NetFlow Monitoring category reports are listed in the following table.

Table F-15 NetFlow Monitoring Category Reports

Report	Description	Drill Down	Parameters
Daily Bandwidth Usage	This report displays a chart and a table to show the bandwidth usage by day.	none	none
Hourly Bandwidth Usage	This report displays a chart and a table to show the bandwidth usage by hour.	none	none
Top Bandwidth Usage by Destination	This report displays a chart and a table to show the bandwidth usage by destination address.	none	none
Top Bandwidth Usage by Destination Port	This report displays a chart and a table to show the bandwidth usage by destination port.	none	none
Top Bandwidth Usage by Source	This report displays a chart and a table to show the bandwidth usage by source address.	none	none

Network Monitoring

This is a sub-category of the Foundation category, focusing on network bandwidth and status events.

The Network Monitoring category is located under the following path.

Foundation\Network Monitoring

The Network Monitoring category reports are listed in the following table.

Table F-16 Network Monitoring Category Reports

Report	Description	Drill Down	Parameters
Top VPN Accesses by User	This report displays a table showing the source user ID and name, and the count of events for VPN access, authorization or authentication events.	none	none

Report	Description	Drill Down	Parameters
Top VPN Event Destinations	This report displays a table showing the VPN destination zone, address and host name, and the count of events for that host, reported by the VPN device, excluding modification events.	none	none
Top VPN Events	This report displays a table showing the VPN event name, source zone and address, destination zone and address, and the count of events for that event reported by the VPN device, excluding modification events.	none	none
Top VPN Event Sources	This report displays a table showing the VPN source zone, address and host name, and the count of events for that source, reported by the VPN device, excluding modification events.	none	none
Traffic Statistics	This report displays two charts and a table. The first chart shows the bytes in and out by hour. The second chart shows the bytes in and out by device. The table shows the hour, firewall zone and address, the transport protocol and the bytes in and out.	none	none
VPN Connection Attempts	This report displays a table showing the source hostname, source user name, destination zone, address and host name, destination user ID and user name and the count of events where the VPN access, authorization or authentication event did not result in failure.	none	none

Report	Description	Drill Down	Parameters
VPN Connection Failures	This report displays a table showing the VPN device zone, address and host name, the VPN event, the source user ID, host name and user name, the destination zone, address, host name and user name, and the count of each event, where the VPN device reports and access, authorization or authentication failure.	none	none

SANS Top 5

This category covers the SANS Top 5 Essential Log Reports (<http://www.sans.org/security-resources/top5-logreports.pdf>). Each of the sub-categories addresses one of the 5 areas.

The following categories are located under the SANS Top 5 category:

- "1 - Attempts to Gain Access through Existing Accounts" on page 619
- "2 - Failed File or Resource Access Attempts" on page 620
- "3 - Unauthorized Changes to Users Groups and Services" on page 621
- "4 - Systems Most Vulnerable to Attack" on page 622
- "5 - Suspicious or Unauthorized Network Traffic Patterns" on page 623

1 - Attempts to Gain Access through Existing Accounts

This is a sub-category of the SANS Top 5 Essential Log Reports. It addresses attempts to gain access to a system through existing accounts.

The 1 - Attempts to Gain Access through Existing Accounts category is located under the following path.

SANS Top 5\1 - Attempts to Gain Access through Existing Accounts

The 1 - Attempts to Gain Access through Existing Accounts category reports are listed in the following table.

Table F-17 1 - Attempts to Gain Access through Existing Accounts Category Reports

Report	Description	Drill Down	Parameters
Number of Failed Logins	This report, based on the SANS Top 5 Essential Log Reports, section 1 - Attempts to Gain Access Through Existing Accounts, displays a table showing the number of failed logins for each hour covered by the report time-range.	none	none

Report	Description	Drill Down	Parameters
Top Users with Failed Logins	This report, based on the SANS Top 5 Essential Log Reports, section 1 - Attempts to Gain Access Through Existing Accounts, displays a table showing the user ID and name, the time and the number of attempts to login to a system during that minute.	none	none

2 - Failed File or Resource Access Attempts

This is a sub-category of the SANS Top 5 Essential Log Reports. It addresses failed file or resource access attempts.

The 2 - Failed File or Resource Access Attempts category is located under the following path.

SANS Top 5\2 - Failed File or Resource Access Attempts

The 2 - Failed File or Resource Access Attempts category reports are listed in the following table.

Table F-18 2 - Failed File or Resource Access Attempts Category Reports

Report	Description	Drill Down	Parameters
Failed Resource Access by Users	This report, based on the SANS Top 5 Essential Log Reports, section 2 - Failed File or Resource Access Attempts, displays a table showing the Resource Type, the User ID and Name, Destination zone and address, the Access Event and number of failed attempts to access a given resource. Clicking on an entry in the Resource Type column will bring up the SANS Top 5 -2- Failed Resource Access by Users Drilldown report, which will show all related events for that resource type.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Resource Type field drill downs to the following report: “Failed Resource Access by Users” on page 620 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “categoryObjectParameter” on page 627

Report	Description	Drill Down	Parameters
Failed Resource Access Events	This report, based on the SANS Top 5 Essential Log Reports, section 2 - Failed File or Resource Access Attempts, displays a table showing the Resource Type, the User ID and Name, Destination zone and address, the Access Event and number of failed attempts to access a given resource. Clicking on an entry in the Resource Type column will bring up the SANS Top 5 -2- Failed Resource Access Events Drilldown report, which will show all related events for that resource type.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> Resource Type field drill downs to the following report: “Failed Resource Access Events” on page 621 <p>This report drills down to itself.</p>	<p>This report prompts you to supply a value for the following parameter.</p> <ul style="list-style-type: none"> “categoryObjectParameter” on page 627

3 - Unauthorized Changes to Users Groups and Services

This is a sub-category of the SANS Top 5 Essential Log Reports. It addresses unauthorized changes to users, groups and services.

The 3 - Unauthorized Changes to Users Groups and Services category is located under the following path.

SANS Top 5\3 - Unauthorized Changes to Users Groups and Services

The 3 - Unauthorized Changes to Users Groups and Services category reports are listed in the following table.

Table F-19 3 - Unauthorized Changes to Users Groups and Services Category Reports

Report	Description	Drill Down	Parameters
Account Modifications	This custom report, based on the SANS Top 5 Essential Log Reports, section 3 - Unauthorized Changes to Users, Groups and Services, displays a chart and a table. The chart shows the top user account modifications. The table shows the source user name, source zone and address, destination user name, destination zone and address, the modification event, and the date of the modification.	none	none

Report	Description	Drill Down	Parameters
Password Changes	This report, based on the SANS Top 5 Essential Log Reports, section 3 - Unauthorized Changes to Users, Groups and Services, displays a table showing the user name, source zone and address, destination zone and address, and the date of password change events.	none	none
User Account Creations	This report, based on the SANS Top 5 Essential Log Reports, section 3 - Unauthorized Changes to Users, Groups and Services, displays a table showing the source user name, the source zone and address, the destination user name, the destination zone and address, the modification event name and the date of the account creation.	none	none
User Account Deletions	This report, based on the SANS Top 5 Essential Log Reports, section 3 - Unauthorized Changes to Users, Groups and Services, displays a table showing the source user name, the source zone and address, the destination user name, the destination zone and address, and the time when a user account was deleted.	none	none
User Account Modifications	This report, based on the SANS Top 5 Essential Log Reports, section 3 - Unauthorized Changes to Users, Groups and Services, displays a table showing the source user name, the source zone and address, the destination user name, the destination zone and address, the modification event name, and the date of the account modification.	none	none

4 - Systems Most Vulnerable to Attack

This is a sub-category of the SANS Top 5 Essential Log Reports. It addresses the systems that are most vulnerable to attack.

The 4 - Systems Most Vulnerable to Attack category is located under the following path.

SANS Top 5\4 - Systems Most Vulnerable to Attack

The 4 - Systems Most Vulnerable to Attack category reports are listed in the following table.

Table F-20 4 - Systems Most Vulnerable to Attack Category Reports

Report	Description	Drill Down	Parameters
Vulnerability Scanner Logs by Host	This report, based on the SANS Top 5 Essential Log Reports, section 4 - Systems Most Vulnerable to Attack, displays a table showing the system zone and address, the vulnerability ID and name, and the number of times that vulnerability has been reported for that system.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> arc_destinationAddress field drill downs to the following report: "Vulnerability Scanner Logs by Host" on page 606 <p>This report drills down to itself.</p>	<p>This report prompts you to supply values for the following parameters.</p> <ul style="list-style-type: none"> "IPAddress" on page 627 "zoneParameter" on page 630
Vulnerability Scanner Logs by Vulnerability	This report, based on the SANS Top 5 Essential Log Reports, section 4 - Systems Most Vulnerable to Attack, displays a table showing the vulnerability ID and name, the zone and address, and the number of times that vulnerability has been reported for that system.	<p>This report provides a drilldown from the following field.</p> <ul style="list-style-type: none"> arc_destinationAddress field drill downs to the following report: "Vulnerability Scanner Logs by Host" on page 606 <p>This report drills down to itself.</p>	<p>This report prompts you to supply values for the following parameters.</p> <ul style="list-style-type: none"> "IPAddress" on page 627 "zoneParameter" on page 630

5 - Suspicious or Unauthorized Network Traffic Patterns

This is a sub-category of the SANS Top 5 Essential Log Reports. It addresses suspicious or unauthorized network traffic patterns.

The 5 - Suspicious or Unauthorized Network Traffic Patterns category is located under the following path.

```
SANS Top 5\5 - Suspicious or Unauthorized Network Traffic
Patterns
```

The 5 - Suspicious or Unauthorized Network Traffic Patterns category reports are listed in the following table.

Table F-21 5 - Suspicious or Unauthorized Network Traffic Patterns Category Reports

Report	Description	Drill Down	Parameters
Alerts from IDS	This report, based on the SANS Top 5 Essential Log Reports, section 5 - Suspicious or Unauthorized Network Traffic Patterns, displays a table showing the device vendor and product, the device event ID, the IDS signature name and the number of times that signature was reported.	none	none
IDS Signature Destinations	This report, based on the SANS Top 5 Essential Log Reports, section 5 - Suspicious or Unauthorized Network Traffic Patterns, displays a table showing the destination zone and address, the device vendor and product and the count of events reported for the address by the IDS.	none	none
IDS Signature Sources	This report, based on the SANS Top 5 Essential Log Reports, section 5 - Suspicious or Unauthorized Network Traffic Patterns, displays a table showing the destination zone and address, the device vendor and product, and the count of each event.	none	none
Top 10 Talkers	This report, based on the SANS Top 5 Essential Log Reports, section 5 - Suspicious or Unauthorized Network Traffic Patterns, displays a table showing the source zone and address, and the number of events coming from each address.	none	none

Report	Description	Drill Down	Parameters
Top 10 Types of Traffic	This report, based on the SANS Top 5 Essential Log Reports, section 5 - Suspicious or Unauthorized Network Traffic Patterns, breaks down the traffic by the Application Protocol, Port number and Transport Protocol, where at least one of the three must be available and the bytes in or bytes out are available. The count is based on the number of base events, presuming that each event with these conditions represents a packet of some type.	none	none
Top Alerts from IDS	This report, based on the SANS Top 5 Essential Log Reports, section 5 - Suspicious or Unauthorized Network Traffic Patterns, displays a chart and a table. The chart shows the top 10 alerts from IDSes. The table shows the Signature ID, the signature name, the device vendor and the number of times that signature was reported.	none	none
Top Destination IPs	This report, based on the SANS Top 5 Essential Log Reports, section 5 - Suspicious or Unauthorized Network Traffic Patterns, displays a table showing the target zone and address and the count of events for each destination address.	none	none
Top IDS Signature Destinations	This report, based on the SANS Top 5 Essential Log Reports, section 5 - Suspicious or Unauthorized Network Traffic Patterns, displays a chart and a table. The chart shows the top signature destinations by address. The table shows the destination zone and address, the device vendor and product, and the count of events to that host.	none	none

Report	Description	Drill Down	Parameters
Top IDS Signature Sources	This report, based on the SANS Top 5 Essential Log Reports, section 5 - Suspicious or Unauthorized Network Traffic Patterns, displays a chart and a table. The chart shows the top signature sources by address. The table shows the source zone and address, the device vendor and product, and the count of events by that host.	none	none
Top Target IPs	This report, based on the SANS Top 5 Essential Log Reports, section 5 - Suspicious or Unauthorized Network Traffic Patterns, displays a table showing the target zone and address, and the number of IDS event reported for that address where the category significance is Compromise or Hostile.	none	none

Parameters

Some reports invoke queries that prompt for field values during report runtime. The values entered for these fields are passed to the query using parameters. To change the value of the parameter that is passed to the query, you can enter a new value when prompted by the report during runtime or you can change the default value of the parameter. SQL wildcards are supported values for parameters; for example, the % wildcard character matches one or more characters. For more information about parameters, see [“Working with Parameters” on page 223](#).

Logger reports invoke queries that use the following parameters:

- [“IPAddress” on page 627](#)
- [“categoryObjectParameter” on page 627](#)
- [“commonlyBlockedPorts” on page 627](#)
- [“destinationAddress” on page 628](#)
- [“destinationPort” on page 628](#)
- [“deviceGroupParameter” on page 628](#)
- [“deviceProduct” on page 628](#)
- [“deviceSeverityParameter” on page 628](#)
- [“deviceVendor” on page 629](#)
- [“dmBandwidthParameter” on page 629](#)
- [“dmConfigurationParameter” on page 629](#)
- [“dmLoginParameter” on page 629](#)
- [“eventNameParameter” on page 630](#)

- [“resourceTypeParameter” on page 630](#)
- [“webPorts” on page 630](#)
- [“zoneParameter” on page 630](#)
- [“zones” on page 631](#)

IPAddress

When a report invokes a query that expects the `IPAddress` parameter as input, the IP Address prompt is displayed during report runtime with a default value of `%`. This is a single value character type (CHAR) parameter that takes an IP address, such as `192.168.35.5`.

This parameter is used with the `LIKE` keyword in the `WHERE` clause of an SQL query. See the Foundation \ Intrusion Monitoring \ Attackers \ Top Attacker Details query object for an example of a query using this parameter.

This parameter is used by the following reports: `Top Attacker Details`, `Top Targets`, `Vulnerability Scanner Logs by Host` and `Vulnerability Scanner Logs by Vulnerability`.

categoryObjectParameter

When a report invokes a query that expects the `categoryObjectParameter` parameter as input, the Resource Type prompt is displayed during report runtime with a default value of

`' /Host/Application/Database', '/Host/Application/Database/Data', '/Host/Application/Service/Email', '/Host/Resource/File'`. This is a multiple value character type (CHAR) parameter that allows the selection of one or more Category Object URIs, such as `Host/Application/Database`.

This parameter is used with the `IN` keyword in the `WHERE` clause of an SQL query. See the Foundation \ Intrusion Monitoring \ Resource Access \ Access Events by Resource query object for an example of a query using this parameter.

This parameter is used by the following reports: `Access Events by Resource`, `Resource Access by Users - Failures`, `Resource Access by Users - Successes-Attempts`, `Failed Resource Access by Users` and `Failed Resource Access Events`.

commonlyBlockedPorts

When a report invokes a query that expects the `commonlyBlockedPorts` parameter as input, the Blocked Ports prompt is displayed during report runtime with all default values listed in the Combo Source panel. This is a multiple value number type (NUMBER) parameter that allows the selection of one or more listed port numbers, such as `135,139`.

This parameter is used with the `IN` keyword in the `WHERE` clause of an SQL query. See the Foundation \ Intrusion Monitoring \ Probes on Blocked Ports query object for an example of a query using this parameter.

This parameter is used by the following report: `Probes on Blocked Ports by Source`.

destinationAddress

When a report invokes a query that expects the `destinationAddress` parameter as input, the Destination IP Address prompt is displayed during report runtime with a default value of `%`. This is a single value character type (CHAR) parameter that takes an IP address, such as `192.168.35.5`.

This parameter is used with the LIKE keyword in the WHERE clause of an SQL query. See the Foundation \ Intrusion Monitoring \ Attackers \ Source Counts by Destination query object for an example of a query using this parameter.

This parameter is used by the following reports: `Source Counts by Destination` and `Attack Events by Destination`.

destinationPort

When a report invokes a query that expects the `destinationPort` parameter as input, the Destination Port prompt is displayed during report runtime with a default value of `80`. This is a single value number type (NUMBER) parameter that allows the entry of one port number, such as `80`.

deviceGroupParameter

When a report invokes a query that expects the `deviceGroupParameter` parameter as input, the Category Device Group prompt is displayed during report runtime with a default value of

`'/Firewall', '/IDS', '/IDS/Host', '/IDS/Host/Antivirus', '/IDS/Host/FileIntegrity', '/IDS/Network', '/IDS/Network/TrafficAnalysis', '/NetworkEquipment', '/NetworkEquipment/Router', '/NetworkEquipment/Switches', '/VPN'`. This is a multiple value character type (CHAR) parameter that allows the selection of one or more Category Object URIs, such as `Host/Application/Database`.

This parameter is used with the IN keyword in the WHERE clause of an SQL query. See the Foundation \ Configuration Monitoring \ Device Configuration Changes query object for an example of a query using this parameter.

This parameter is used by the following reports: `Device Configuration Changes`, `Device Configuration Events` and `Device Misconfigurations`.

deviceProduct

When a report invokes a query that expects the `deviceProduct` parameter as input, the Device Product prompt is displayed during report runtime with a default value of `%`. This is a single value character type (CHAR) parameter that takes a string, such as `Snort`.

This parameter is used with the LIKE keyword in the WHERE clause of an SQL query. See the Foundation \ Configuration Monitoring \ Vulnerability Scanner Logs by Host query object for an example of a query using this parameter.

This parameter is used by the following reports: `Vulnerability Scanner Logs by Host` and `Vulnerability Scanner Logs by Vulnerability`.

deviceSeverityParameter

When a report invokes a query that expects the `deviceSeverityParameter` parameter as input, the Device Severity prompt is displayed during report runtime with a default value

of %. This is a single value character type (CHAR) parameter that takes a string, such as High.

This parameter is used with the LIKE keyword in the WHERE clause of an SQL query. See the Foundation \ Intrusion Monitoring \ Attackers \ Source Counts by Device Severity query object for an example of a query using this parameter.

This parameter is used by the following reports: Source Counts by Device Severity and Destination Counts by Device Severity.

deviceVendor

When a report invokes a query that expects the deviceVendor parameter as input, the Device Vendor prompt is displayed during report runtime with a default value of %. This is a single value character type (CHAR) parameter that takes a string, such as Snort.

This parameter is used with the LIKE keyword in the WHERE clause of an SQL query. See the Foundation \ Configuration Monitoring \ Vulnerability Scanner Logs by Host query object for an example of a query using this parameter.

This parameter is used by the following reports: Vulnerability Scanner Logs by Host and Vulnerability Scanner Logs by Vulnerability.

dmBandwidthParameter

When a report invokes a query that expects the dmBandwidthParameter parameter as input, the Device Type prompt is displayed during report runtime with a default value of all. This is a single value character type (CHAR) parameter that allows the selection of predefined value, such as Firewall.

This parameter is used with the LIKE keyword in the WHERE clause of an SQL query. See the Foundation \ Configuration Monitoring \ Vulnerability Scanner Logs by Host query object for an example of a query using this parameter.

This parameter is used by the following reports: Bandwidth Usage by Hour, Bandwidth Usage by Protocol, Top Bandwidth Hosts and Top Hosts by Number of Connections.

dmConfigurationParameter

When a report invokes a query that expects the dmConfigurationParameter parameter as input, the Device Type prompt is displayed during report runtime with a default value of all. This is a single value character type (CHAR) parameter that allows the selection of predefined value, such as Firewall.

This parameter is used with the LIKE keyword in the WHERE clause of an SQL query.

This parameter is used by the following reports: Configuration Changes by Type and Configuration Changes by User.

dmLoginParameter

When a report invokes a query that expects the dmLoginParameter parameter as input, the Device Type prompt is displayed during report runtime with a default value of all. This is a single value character type (CHAR) parameter that allows the selection of predefined value, such as Firewall.

This parameter is used with the LIKE keyword in the WHERE clause of an SQL query. See the Device Monitoring \ CrossDevice \ Failed Login Attempts query object for an example of a query using this parameter.

This parameter is used by the following reports: Failed Login Attempts, Failed Logins by Destination Address, Failed Logins by Source Address, Failed Logins by User, Login Event Audit, Password Changes, Successful Logins by Destination Address, Successful Logins by Source Address and Successful Logins by User.

eventNameParameter

When a report invokes a query that expects the `eventNameParameter` parameter as input, the Event Name prompt is displayed during report runtime with a default value of %. This is a single value character type (CHAR) parameter that takes a string, such as Connector Raw Event Statistics.

This parameter is used with the LIKE keyword in the WHERE clause of an SQL query. See the Foundation \ Intrusion Monitoring \ Targets \ Destination Counts by Event Name query object for an example of a query using this parameter.

This parameter is used by the following report: Destination Counts by Event Name.

resourceTypeParameter

When a report invokes a query that expects the `resourceTypeParameter` parameter as input, the Resource Type prompt is displayed during report runtime with a default value of /Host/Application/Database. This is a single value character type (CHAR) parameter that takes a string, such as /Host/Application/Database.

This parameter is used with the LIKE keyword in the WHERE clause of an SQL query.

webPorts

When a report invokes a query that expects the `webPorts` parameter as input, the Web Ports prompt is displayed during report runtime with all default values listed in the Combo Source panel. This is a multiple value number type (NUMBER) parameter that allows the selection of one or more listed port numbers, such as 80,443.

This parameter is used with the IN keyword in the WHERE clause of an SQL query. See the Foundation \ Intrusion Monitoring \ Top Web Traffic query object for an example of a query using this parameter.

This parameter is used by the following reports: Top Web Traffic and Top Machines Accessing the Web.

zoneParameter

When a report invokes a query that expects the `zoneParameter` parameter as input, the Zone prompt is displayed during report runtime with a default value of %. This is a single value character type (CHAR) parameter that takes an IP address, such as /All Zones/ArcSight System/Private Address Space Zones/RFC1918: 10.0.0.0-10.255.255.255.

This parameter is used with the LIKE keyword in the WHERE clause of an SQL query. See the Foundation \ Intrusion Monitoring \ Attackers \ Top Attacker Details query object for an example of a query using this parameter.

This parameter is used by the following reports: Top Attacker Details, Top Attackers, Top Targets, Vulnerability Scanner Logs by Host and Vulnerability Scanner Logs by Vulnerability.

zones

When a report invokes a query that expects the zones parameter as input, the Zone prompt is displayed during report runtime with a default value of %. This is a multiple value character type (CHAR) parameter that allows the selection of one or more Category Object URIs, such as /All Zones/ArcSight System/Private Address Space Zones/RFC1918: 10.0.0.0-10.255.255.255,/All Zones/ArcSight System/Private Address Space Zones/RFC1918: 192.168.0.0-192.168.255.255.

This parameter is used with the IN keyword in the WHERE clause of an SQL query. See the Foundation \ Intrusion Monitoring \ Attackers \ Source Counts by Destination query object for an example of a query using this parameter.

This parameter is used by the following reports: Source Counts by Destination and Attack Events by Destination.

System Filters

Logger provides the system filters listed in the following table:

Table F-22 System Filters

Filter	Type	Description
Configuration - Configuration Changes (Unified)	Unified Query	This filter looks for events categorized as configuration changes events.
Configuration - System Configuration Changes (CEF format)	Regular Expression	This filter looks for events categorized as configuration changes events. It is a Regular Expression filter and can be used to create alerts.
Events - CEF	Regular Expression	This filter looks for all CEF formatted events. It is a Regular Expression filter and can be used to create alerts.
Events - Event Counts by Destination	Unified Query	This filter looks for all CEF events that have a destination address and shows a chart.
Events - Event Counts by Source	Unified Query	This filter looks for all CEF events that have a source address and shows a chart.
Events - High and Very High Severity CEF Events	Regular Expression	This filter looks for CEF events with a high or very high severity. It is a Regular Expression filter and can be used to create alerts.
Events - High and Very High Severity Events (Unified)	Unified Query	This filter looks for CEF events with a high or very high severity.

Filter	Type	Description
Firewall - Deny	Unified Query	This filter looks for events with deny or shun.
Firewall - Drop	Unified Query	This filter looks for drop events that are not database related.
Firewall - Permit	Unified Query	This filter looks for events that have the word permit.
Intrusion - Malicious Code (CEF format)	Regular Expression	This filter looks for CEF events categorized to indicate malicious code. It is a Regular Expression filter and can be used to create alerts.
Intrusion - Malicious Code (Unified)	Unified Query	This filter looks for CEF events categorized to indicate malicious code.
Logins - All Logins (CEF format)	Regular Expression	This filter looks for CEF events categorized as authentication events. It is a Regular Expression filter and can be used to create alerts.
Logins - All Logins (Non-CEF format)	Regular Expression	This filter looks for non-CEF format events with words indicating it is an authentication event. It is a Regular Expression filter and can be used to create alerts.
Logins - All Logins (Unified)	Unified Query	This filter looks for CEF events categorized as authentication events.
Logins - Failed Logins	Unified Query	This filter looks for failure events related to logins, user authentication and user authorization.
Logins - Successful Logins (Non-CEF format)	Regular Expression	This filter looks for events with keywords indicating a successful login attempt. It is a Regular Expression filter and can be used to create alerts.
Logins - Successful Logins (CEF format)	Regular Expression	This filter looks for CEF events categorized as successful login events. It is a Regular Expression filter and can be used to create alerts.
Logins - Successful Logins (Unified)	Unified Query	This filter looks for CEF events categorized as successful login events.
Logins - Unsuccessful Logins (Non-CEF format)	Regular Expression	This filter looks for failure events related to logins, user authentication and user authorization. It is a Regular Expression filter and can be used to create alerts.
Logins - Unsuccessful Logins (CEF format)	Regular Expression	This filter looks for failure events related to logins, user authentication and user authorization. It is a Regular Expression filter and can be used to create alerts.
Logins - Unsuccessful Logins (Unified)	Unified Query	This filter looks for failure events categorized as login events.
Network - DHCP Lease Events	Unified Query	This filter looks for DHCP lease related events.
Network - Port Links Up and Down	Unified Query	This filter looks for port or link status messages.
Network - Protocol Links Up and Down	Unified Query	This filter looks for protocol status messages.
SystemAlert - CPU Utilization Above 90% (CEF format)	Regular Expression	This filter looks for internal events indicating that CPU utilization is above 90%. It is a Regular Expression filter and can be used to create alerts.

Filter	Type	Description
SystemAlert - CPU Utilization Above 90% (Unified)	Unified Query	This filter looks for internal events indicating that CPU utilization is above 90%.
SystemAlert - CPU Utilization Above 95% (CEF format)	Regular Expression	This filter looks for internal events indicating that CPU utilization is above 95%. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - CPU Utilization Above 95% (Unified)	Unified Query	This filter looks for internal events indicating that CPU utilization is above 95%.
SystemAlert - Device Configuration Changes (CEF format)	Regular Expression	This filter looks for internal events indicating a Logger configuration change. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Device Configuration Changes (Unified)	Unified Query	This filter looks for internal events indicating a Logger configuration change.
SystemAlert - Filter Configuration Changes (CEF format)	Regular Expression	This filter looks for internal events indicating a Logger filter change. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Filter Configuration Changes (Unified)	Unified Query	This filter looks for internal events indicating a Logger filter change.
SystemAlert - High CPU Temperature (CEF format)	Regular Expression	This filter looks for internal events indicating potential CPU over-heating. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - High CPU Temperature (Unified)	Unified Query	This filter looks for internal events indicating potential CPU over-heating.
SystemAlert - Bad Fan (CEF format)	Regular Expression	This filter looks for Logger appliance internal events related to fan failure. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Power Supply Failure (CEF format)	Regular Expression	This filter looks for internal events indicating that a power supply has failed. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Power Supply Failure (Unified)	Unified Query	This filter looks for internal events indicating that a power supply has failed.
SystemAlert - RAID Status Battery Failure (CEF format)	Regular Expression	This filter looks for internal events indicating that the RAID battery backup unit (BBU) has failed. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - RAID Status Battery Failure (Unified)	Unified Query	This filter looks for internal events indicating that the RAID battery backup unit (BBU) has failed.
SystemAlert - Disk Failure (CEF format)	Regular Expression	This filter looks for Logger appliance internal events indicating a disk failure. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Disk Failure (Unified)	Unified Query	This filter looks for Logger appliance internal events indicating a disk failure.
SystemAlert - RAID Controller Issue (CEF format)	Regular Expression	This filter looks for internal events indicating that a RAID disk has failed. It is a Regular Expression filter and can be used to create alerts.

Filter	Type	Description
SystemAlert - RAID Controller Issue (Unified)	Unified Query	This filter looks for internal events indicating that a RAID disk has failed.
SystemAlert - Root Partition Free Space Below 5% (Unified)	Unified Query	This filter looks for internal events indicating that the root partition disk free space is below 5%.
SystemAlert - Root Partition Free Space Below 10% (Unified)	Unified Query	This filter looks for internal events indicating that the root partition disk free space is below 10%.
SystemAlert - Root Partition Free Space Below 10% (CEF format)	Regular Expression	This filter looks for internal events indicating that the root partition disk free space is below 10%. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Root Partition Free Space Below 5% (CEF format)	Regular Expression	This filter looks for internal events indicating that the root partition disk free space is below 5%. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Storage Configuration Changes (CEF format)	Regular Expression	This filter looks for Logger internal events related to changes of the storage configuration. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Storage Configuration Changes (Unified)	Unified Query	This filter looks for Logger internal events related to changes of the storage configuration.
SystemAlert - Storage Group Usage Above 90% (CEF format)	Regular Expression	This filter looks for Logger internal events indicating that the storage group usage is above 90%. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Storage Group Usage Above 90% (Unified)	Unified Query	This filter looks for Logger internal events indicating that the storage group usage is above 90%.
SystemAlert - Storage Group Usage Above 95% (CEF format)	Regular Expression	This filter looks for Logger internal events indicating that the storage group usage is above 95%. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Storage Group Usage Above 95% (Unified)	Unified Query	This filter looks for Logger internal events indicating that the storage group usage is above 95%.
SystemAlert - Zero Events Incoming (CEF format)	Regular Expression	This filter looks for Logger internal events indicating that the no events are being received by Logger. It is a Regular Expression filter and can be used to create alerts.
SystemAlert - Zero Events Incoming (Unified)	Unified Query	This filter looks for Logger internal events indicating that the no events are being received by Logger.
SystemAlert - Zero Events Outgoing (CEF format)	Regular Expression	This filter looks for Logger internal events indicating that the no events are being forwarded by Logger. It is a Regular Expression filter and can be used to create alerts.

Filter	Type	Description
SystemAlert - Zero Events Outgoing (Unified)	Unified Query	This filter looks for Logger internal events indicating that the no events are being forwarded by Logger.
SystemStatus - CPU Utilization by Connector Host	Unified Query	This filter looks for SmartConnector system health events and charts the CPU utilization by host.
SystemStatus - Disk Utilization by Connector Host	Unified Query	This filter looks for SmartConnector system health events and charts the disk utilization by host.
SystemStatus - Memory Utilization by Connector Host	Unified Query	This filter looks for SmartConnector system health events and charts the memory utilization by host.
Unix - CRON related events	Unified Query	This filter looks for events with the cron keyword.
Unix - IO Errors and Warnings	Unified Query	This filter looks for I/O events with error or warning keywords.
Unix - PAM and Sudo Messages	Unified Query	This filter looks for events with the keywords PAM or sudo.
Unix - Password Changes	Unified Query	This filter looks for events related to password changes.
Unix - SAMBA Events	Unified Query	This filter looks for events related to SAMBA.
Unix - SSH Authentications	Unified Query	This filter looks for SSH authentication events.
Unix - User and Group Additions	Unified Query	This filter looks for events related to adding users or groups.
Unix - User and Group Deletions	Unified Query	This filter looks for events related to deleting users or groups.
Windows - Account Added to Global Group	Unified Query	This filter looks for non-CEF events related to adding a Windows account to a Global Group.
Windows - Account Added to Global Group (CEF)	Unified Query	This filter looks for CEF events related to adding a Windows account to a Global Group.
Windows - Audit Policy Change	Unified Query	This filter looks for non-CEF events related to Windows Audit Policy changes.
Windows - Audit Policy Change (CEF)	Unified Query	This filter looks for CEF events related to Windows Audit Policy changes.
Windows - Change Password Attempt	Unified Query	This filter looks for non-CEF events related to Windows password changes.
Windows - Change Password Attempt (CEF)	Unified Query	This filter looks for CEF events related to Windows password changes.
Windows - Global Group Created	Unified Query	This filter looks for non-CEF events related to the creation of Windows global groups
Windows - Global Group Created (CEF)	Unified Query	This filter looks for CEF events related to the creation of Windows global groups.

Filter	Type	Description
Windows - Logon Bad User Name or Password	Unified Query	This filter looks for non-CEF events related to Windows logon failures.
Windows - Logon Bad User Name or Password (CEF)	Unified Query	This filter looks for CEF events related to Windows logon failures.
Windows - Logon Local User	Unified Query	This filter looks for non-CEF events related to Windows logons to the local system.
Windows - Logon Local User (CEF)	Unified Query	This filter looks for CEF events related to Windows logons to the local system.
Windows - Logon Remote User	Unified Query	This filter looks for CEF events related to Windows logons to a remote system.
Windows - Logon Remote User (CEF)	Unified Query	This filter looks for CEF events related to Windows logons to a remote system.
Windows - Logon Unexpected Failure	Unified Query	This filter looks for non-CEF events related to Windows logons with an unexpected failure.
Windows - Logon Unexpected Failure (CEF)	Unified Query	This filter looks for CEF events related to Windows logons with an unexpected failure.
Windows - New Process Creation	Unified Query	This filter looks for non-CEF events related to the creation of new Windows processes.
Windows - New Process Creation (CEF)	Unified Query	This filter looks for CEF events related to the creation of new Windows processes.
Windows - Pre-Authentication Failure	Unified Query	This filter looks for non-CEF events related to failures with Windows pre-authentication.
Windows - Pre-Authentication Failure (CEF)	Unified Query	This filter looks for CEF events related to failures with Windows pre-authentication.
Windows - Special Privileges Assigned to New Logon	Unified Query	This filter looks for non-CEF events related to logons for accounts with special privileges (power user or administrator-like accounts).
Windows - Special Privileges Assigned to New Logon (CEF)	Unified Query	This filter looks for CEF events related to logons for accounts with special privileges (power user or administrator-like accounts).
Windows - User Account Changed	Unified Query	This filter looks for non-CEF events related to user account changes.
Windows - User Account Changed (CEF)	Unified Query	This filter looks for CEF events related to user account changes.
Windows - User Account Password Set	Unified Query	This filter looks for non-CEF events related to user account password changes.
Windows - User Account Password Set (CEF)	Unified Query	This filter looks for CEF events related to user account password changes.
Windows - Windows Events (CEF)	Unified Query	This filter looks for all CEF events that are generated by Microsoft Windows.

Destination Runtime Parameters

The following table describes the destination parameters you can configure. The parameters listed in the table are not available for all destinations. The user interface automatically displays the parameters valid for a destination. For step-by-step instructions on updating the runtime parameters of a destination, see [“Editing Destination Parameters” on page 487](#).

Name Fields	Value Fields
Batching	Connectors can batch events to increase performance and optimize network bandwidth. When activated, connectors create blocks of events and send them when they either (1) reach a certain size or (2) the time window expires, whichever occurs first. You can also prioritize batches by severity, forcing the connector 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 (5, 10, 20, 50, 100 , 200, 300 events).
Enable Batching (in seconds)	The connector sends the events if this time window expires (1, 5 , 10, 15, 30, 60).
Batch By	This is Time Based if the connector should send batches as they arrive (the default) or Severity Based if the connector should send batches based on severity (batches of Highest Severity events sent first).
Time Correction	The values you set for these fields establish forward and backward time limits, that if exceeded, cause the connector to automatically correct the time reported by the device.
Use Connector Time as Device Time	Override the time the device reports and instead use the time at which the connector received the event. This option assumes that the connector will be more likely to report the correct time. (No Yes)
Enable Device Time Correction (in seconds)	The connector can adjust the time reported by the device <code>Detect Time</code> , 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. The default is 0 .
Enable Connector Time Correction (in seconds)	The connector can also adjust the time reported by the connector itself, using this setting. This is for informational purposes only and allows you to modify the local time on the connector. This should be a temporary setting. The recommended way to synchronize clocks between Manager and connectors is the NTP protocol. The default is 0 .

Set Device Time Zone To 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 connector 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. Default: **Disabled**.

Device Time Auto-correction

Future Threshold The connector sends the internal alert if the detect time is greater than the connector time by `Past Threshold` seconds.

Past Threshold The connector sends the internal alert 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

These are the time span and frequency factors for doing device-time auto-correction.

Future Threshold The number of seconds by which to extend the connector's forward threshold for time checking. The 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** (3,600 seconds).

Frequency The connector checks its future and past thresholds at intervals specified by this number of seconds. Default is **1 minute** (60 seconds).

Cache

Changing these settings will not affect the events cached, it will only affect new events sent to the cache.

Cache Size Connectors use a compressed disk cache to hold large volumes of events when the ArcSight Manager is down or when the connector 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 connector drops the oldest events to free up disk cache space. (5 MB, 50 MB, 100 MB, 150 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 after the Notification Threshold is reached. (1 minute, 5 minutes, **10 minutes**, 30 minutes, 60 minutes.)

Network

Heartbeat Frequency This setting controls how often the connector sends a heartbeat message to the destination. The default is **10 seconds**, but it can go from **5 seconds** to **10 minutes**.

Note: The heartbeat is also used to communicate with the connector; 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 connector.

Enable Name Resolution The connector tries to resolve IP addresses to hostnames, and hostnames to IP addresses, if required and if the event rate allows. This setting controls this functionality. The Source, Target and Device IP addresses and Hostnames might also be affected by this setting. By default, name resolution is enabled (**Yes**).

Name Resolution Host Name Only	Default: Yes .
Name Resolution Domain From E-mail	Default: Yes .
Clear Host Names Same as IP Addresses	Default: Yes .
Don't Resolve Host Names Matching	NA
Don't Reverse-Resolve IP Ranges	NA
Limit Bandwidth To	A list of bandwidth options you can use to constrain the connector's output over the network. (Disabled , 1 kbit/sec to 100 Mb/s/sec.)
Transport Mode	You can configure the connector to cache to disk all the processed events it receives. This is equivalent to pausing the connector. 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)).
Address-based Zone Population Defaults Enabled	This field applies to 3.0 ArcSight Managers. This field is not relevant in ESM 3.5 because the system has integral zone mapping. Default: Yes .
Address-based Zone Population	This field applies to 3.0 ArcSight Managers. This field is not relevant in ESM 3.5 because the system has integral zone mapping.
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.
Source Zone URI	When populated, this field shows the URI of the zone associated with the connector's source address. This field is present for ESM 3.0 compatibility. It is not relevant in ESM 3.5 because of integral zone mapping.
Source Translated Zone URI	When populated, this field shows the URI of the zone associated with the connector's translated source address. The translation is presumed to be NAT (network address translation). This field is present for ESM 3.0 compatibility. It is not relevant in ESM 3.5 because of integral zone mapping.
Destination Zone URI	When populated, this field shows the URI of the zone associated with the connector's destination address. This field is present for ESM 3.0 compatibility. It is not relevant in ESM 3.5 because of integral zone mapping.
Destination Translated Zone URI	When populated, this field shows the URI of the zone associated with the connector's translated destination address. The translation is presumed to be NAT (network address translation). This field is present for ESM 3.0 compatibility. It is not relevant in ESM 3.5 because of integral zone mapping.
Connector Zone URI	When populated, this field shows the URI of the zone associated with the connector's address. This field is present for ESM 3.0 compatibility. It is not relevant in ESM 3.5 because of integral zone mapping.

Connector Translated Zone URI	When populated, this field shows the URI of the zone associated with the connector's translated address. The translation is presumed to be NAT (network address translation). This field is present for ESM 3.0 compatibility. It is not relevant in ESM 3.5 because of integral zone mapping.
Device Zone URI	When populated, this field shows the URI of the zone associated with the device's address. This field is present for ESM 3.0 compatibility. It is not relevant in ESM 3.5 because of integral zone mapping.
Device Translated Zone URI	When populated, this field shows the URI of the zone associated with the device's translated address. The translation is presumed to be NAT (network address translation). This field is present for ESM 3.0 compatibility. It is not relevant in ESM 3.5 because of integral zone mapping.
Field Based Aggregation	<p>This feature is an extension of basic connector aggregation. Basic aggregation aggregates two events if, and only if, all the fields of the two events are the same (the only difference being the detect time). However, field-based aggregation implements a less strict aggregation mechanism; two events are aggregated if only the selected fields are the same for both alerts. It is important to note that field-based aggregation creates a new alert that contains only the fields that were specified, so the rest of the fields are ignored.</p> <p>Connector 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.</p>
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, or otherwise contained the same selected fields, and you select an event threshold of 100, you will 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	Enter one or more fields, if applicable, to use as the basis for aggregating the events the connector collects. 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-hostnames. Names can contain no spaces and the first letter should not be capitalized.
Fields to Sum	Enter one or more fields, if applicable, to use as the basis for aggregating the events the connector collects.
Preserve Common Fields	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	<p>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).</p> <p>Connector 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.</p>

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, or otherwise contained the same selected fields, and you select an event threshold of 100, you will 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

Preserve Raw Event	For some devices, a raw event 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. 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. The default is No . If you choose Yes , the serialized representation of the "Raw Event" is sent to the destination and preserved in the Raw Event field.
Turbo Mode	<p>If your configuration, reporting, and analytic usage permits, you can greatly accelerate the transfer of a sensor's event information through connectors by choosing one of two "turbo" (narrower data bandwidth) modes. The default transfer mode is called Complete, which passes all the data arriving from the device, including any additional data (custom, or vendor-specific).</p> <p>Complete mode does indeed use all the database performance advances of ArcSight ESM 3.x.</p> <p>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 on reports, rules, threat resolution from a given device before selecting it.</p> <p>The specific event attributes that apply to these modes in your enterprise are defined in the self-documented \$ARCSIGHT_HOME/config/connector/agent.properties file for the ArcSight Manager. Because these properties might have been adjusted for your needs, you should refer to this file for definitive lists. Only scanner connectors need to run in Complete mode, to capture the additional data.</p> <p>Note: Connector Turbo Modes are superseded by the Turbo Mode in use by the ArcSight Managers processing their events. For example, a Manager set to Faster will not pass all the data possible for a connector that is set for the default of Complete.</p>

Enable Aggregation (in seconds)	<p>When enabled, aggregates two or more events on the basis of the selected time value. (Disabled, 1, 2, 3, 4, 5, 10, 30, 60)</p> <p>The aggregation is performed on one or more matches for a fixed subset of fields:</p> <ul style="list-style-type: none">• 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) <p>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.</p>
Limit Event Processing Rate	<p>You can moderate the connector's burden on the CPU by reducing its processing rate. This can also be a means of dealing with the effects of event bursts.</p> <p>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).</p> <p>Note: The effect of this option varies with the category of connector in use, as described in the connector Processing Categories table below.</p>
Fields to Obfuscate	
Store Original Time in	Disabled or Flex Date 1.
Enable Port-Service Mapping	Default: No .
Enable User Name Splitting	Default: No .
Split File Name into Path and Name	Default: No .
Event Integrity Algorithm	Disabled , SHA-1, SHA-256, SHA-512, or MD5.
Generate Unparsed Events	Default: No .

Preserve System Health Events Yes, **No**, or Disabled.

Enable Device Status Monitoring (in minutes) **Disabled** or 1, 2, 3, 4, 5, 10, 30, 60, or 120 minutes.

Filters

Filter Out NA

"Very High Severity" Event Definition NA

"High Severity" Event Definition NA

"Medium Severity" Event Definition NA

"Low Severity" Event Definition NA

"Unknown Severity" Event Definition NA

Payload Sampling (When available.)

Max. Length Discard, 128 bytes, **256 bytes**, 512 bytes, 1 kbyte

Mask Non-Printable Characters Default: **False**.

Appendix H

Restoring Factory Settings

This section describes how to restore your appliance to its original factory settings by over-writing the current files with an image of the original system.



Restoring an appliance to its original factory settings **irrevocably deletes all event data** and some configuration settings.

Before Restoring Your System

You must remember the following cautions and guidelines before you restore to factory settings.

- If your Logger running 5.1 or later has multipath SAN enabled and you encounter one of these situations:
 - ◆ RMA the system to HP and receive a new system that is either running Logger 5.0 Patch 3 or earlier
 - ◆ Factory restore the system to its factory default settings, which resets the Logger version to 5.0 Patch 3 or earlier

To restore your Logger to its last working state—running version 5.1 or later, with multipath enabled—make sure that you upgrade your system to Logger 5.1 or later before attaching the LUN.

- After restoring, you can to restore backups of your data and configuration settings.
- When restoring the configuration of the Logger from a backup, first, ensure that the appliance is restoration, then complete the upgrade to the desired version.

Restoring Your System

Instructions for performing a factory reset vary by the appliance model. Refer to the appropriate section for your appliance.

- [“Restoring the LX500” on page 645](#)
- [“Restoring LX400 and Earlier Appliance Models” on page 647](#)

Restoring the LX500

You can restore LX500 appliances to the original factory settings by using the built-in System Restore utility.

To restore an LX500 appliance:

- 1 Attach a keyboard, monitor, and mouse directly to the appliance or, if your appliance is configured for remote access through iLO, you can use that functionality to access the appliance console. See [“Set Up the Logger Appliance for Remote Access” on page 39](#) for more information.



- 2 Log into the appliance. Type `reboot` at the command prompt, and then press Enter.

As the system reboots, messages scroll by. As soon as a message like the following appears on the screen, press any key on your keyboard.

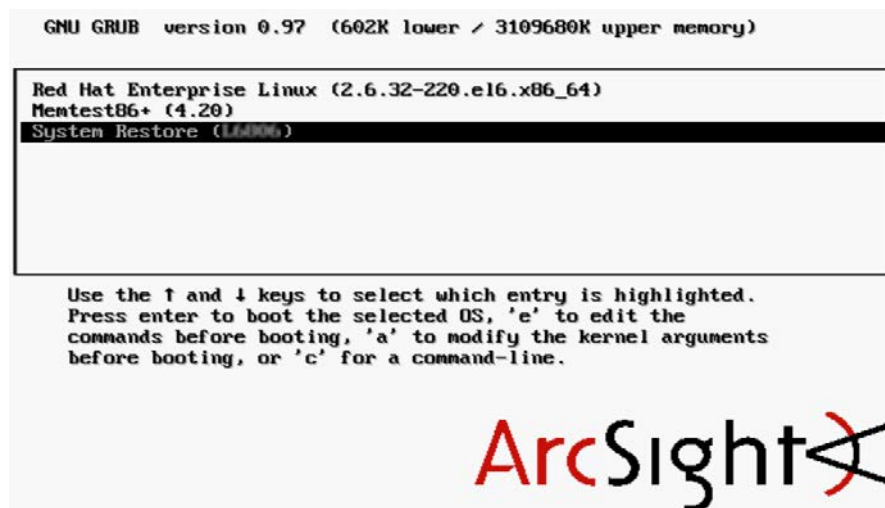
```
Press any key to enter the menu
Booting Red Hat Enterprise Linux <version> in N seconds...
```

This message is displayed for a very short time. Make sure you press a key on your keyboard quickly; otherwise, the appliance will continue to boot normally. You need to press the key after the BIOS is done booting but before the OS boots.

If the OS boots, you will see something like the screenshot below. You'll need to try again in that case.



- 3 The GNU GRUB window opens.



Use the mouse or arrow keys to select **System Restore L<XXXX>** and press **Enter**.

- 4 System Restore automatically detects and displays the archive image. The image is named following the pattern `YYYY-MM-DD_LX500_L<XXXX>.ari`, where `YYYY-MM-DD` is the date, `LX500` is the appliance version and `L<XXXX>` is the appliance build number.
- 5 Press **F1** (AUTOSELECT) to automatically map the Source Image, displayed in the top panel, to the Target Disk, displayed in the bottom panel. The restore image name is displayed in the right-most column.
- 6 Optionally, press **F10** (VERIFY) to check the archive for damage before performing the restore. Once the archive has been verified, press **Enter** to continue.
- 7 Press **F2** (RESTORE) to begin the restore process. A dialog box asks whether you want to restore. Press **y** to proceed with the restore or **n** to cancel.
- 8 Progress bars show the status of the restoration.



Do not interrupt or power-down the appliance during the restore process. Interrupting the restore process may force the system into a state from which it cannot be recovered.


- 9 When the restore process is complete, press **F12** to reboot the appliance. A dialog box asks whether you want to reboot. Press **y** to proceed with the reboot.

Restoring LX400 and Earlier Appliance Models

You can restore LX400 and earlier appliances to the original factory settings by using the built-in Acronis True Image software.

To restore LX400 and earlier appliance models:

- 1 Attach a keyboard, monitor, and mouse directly to the appliance or, if your appliance is configured for remote access through iLO, you can use that functionality to access the appliance console. See [“Set Up the Logger Appliance for Remote Access”](#) on page 39 for more information.



- 2 Log into the appliance. Type `reboot` at the command prompt, and then press **Enter**.

As the system reboots, messages scroll by. As soon as a message like the following appears on the screen, press any key on your keyboard.

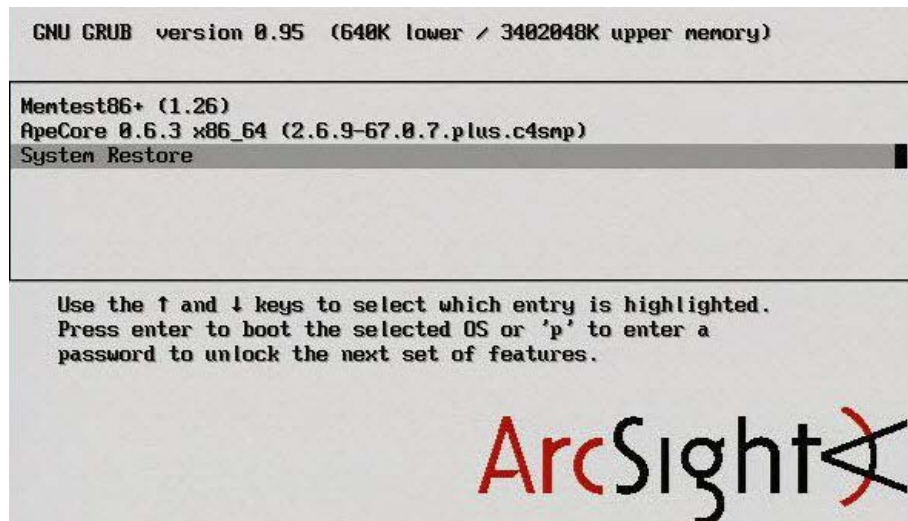
```
Press any key to enter the menu
Booting Red Hat Enterprise Linux <version> in N seconds...
```

This message is displayed for a very short time. Make sure you press a key on your keyboard quickly; otherwise, the appliance will continue to boot normally. You need to press the key after the BIOS is done booting but before the OS boots.

If the OS boots, you will see something like the screenshot below. You'll need to try again in that case.



- 3 The session viewer opens.



Use the mouse or arrow keys to select **System Restore** and press **Enter**.

- 4 In the **Acronis True Image Echo Server** dialog box, select **Recovery** from the **Pick a Task** list and then press Enter.
- 5 When the Restore Data Wizard starts, click **Next** to continue.
- 6 On the Welcome to the Restore Data Wizard page, click **Next** to continue.
- 7 On the **Backup Archive Selection** page, select **Acronis Secure Zone** and click **Next**. You will have a chance to review the choices you make on this page and the wizard pages that follow before initiating the restore process.
- 8 On the **Restoration Type Selection** page, select **Restore disks or partitions** and click **Next**. Only choose other options if specifically directed to do so by customer support.
- 9 On the **Partition or Disk to Restore** page, select the entire drive, labeled **cciss/c0d0** or **sda** (depending on the appliance model) and click **Next**.
- 10 On the **NT Signature selection for image restoration** page, select how you want the NT signature for the restored disk to be processed and click **Next**.
- 11 On the **Restored Hard disk Location** page, select the drive to restore (**cciss/c0d0** or **sda**) and click **Next**.
- 12 On the **Non-empty Destination Hard Disk Drive** page, select **Yes, I want to delete all partitions on the destination hard disk drive before restoring** and click **Next**.

- 13** On the **Next Selection** page, select **No, I do not** and click **Next** (there are no other partitions or disks to restore).
- 14** On the **Restoration Options** page, select **Validate backup archive for the data restoration process** if you want to validate the archive before resetting the appliance. Select **Reboot the computer automatically after the restoration is finished** if you want to reboot the appliance automatically. Click **Next**.
- 15** Review the checklist of operations to be performed and click **Proceed** to begin factory reset. Click **Back** to revisit previous pages.



Do not interrupt or power-down the Logger appliance during the restore process. Interrupting the restore process may force the system into a state from which it cannot be recovered.

Progress bars show the status of the current operation and the total progress.

- 16** When you see a message indicating that the data was restored successfully, click **OK**.

If you specified automatic reboot, the appliance reboots when the restore is complete. Otherwise, reboot manually.

Appendix I

Logger Search From an ArcSight Console

If you have ArcSight Logger and ArcSight ESM deployed in your network infrastructure, you can perform a Logger search operation directly from your ArcSight Console. This appendix discusses how use the integrated search functionality. It includes information on the following topics.

[“Understanding the Integrated Search Functionality” on page 651](#)

[“Prerequisites” on page 652](#)

[“Setup and Configuration” on page 653](#)

[“Supported Search Options” on page 654](#)

[“Guidelines” on page 654](#)

[“Searching on Logger From ArcSight Console” on page 655](#)

Understanding the Integrated Search Functionality

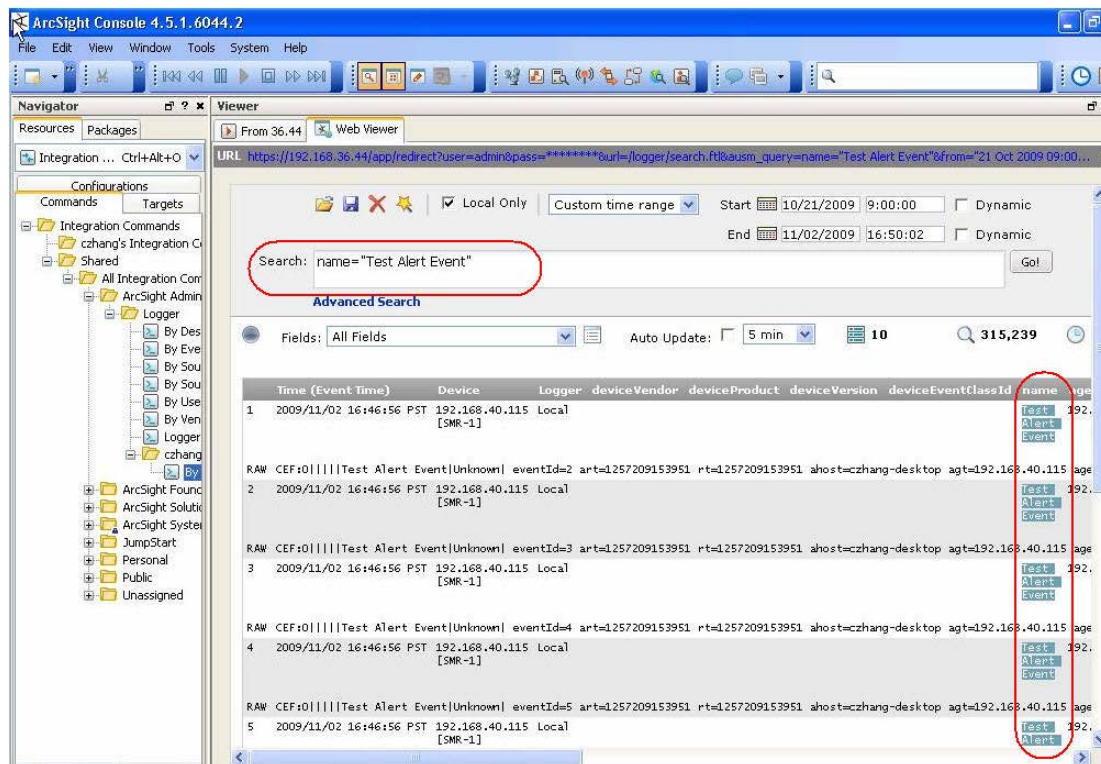
There are two ways to perform a search operation on Logger from an ArcSight Console:

- Search — a regular search operation in which you can specify search options.
- Quick search — a search operation based on field and value you select in an ArcSight Console active channel; you are not prompted for any search options.

To run a Logger search, right click on an event in an active channel of the ArcSight Console to display a menu to select a search method — Logger Search or Logger Quick Search.

If you select Logger Search, you need to select search options such as Event Name, Destination, Source, and so on, and the Logger appliances on which the search should be run (if there are multiple Logger appliances). If you select Logger Quick Search, you are not prompted to specify any search options because the search is run on the field you had clicked on.

The search results are displayed in the ArcSight Console, as shown in the following figure:



Before you can run a search operation on Logger from ArcSight Console, you need to set up parameters in the ArcSight Console that are used to authenticate the user who performs the search. Prior to Logger 5.1, only basic authentication (user name and password) was available; however, starting with Logger 5.1, a One Time Password (OTP) option is available. This option makes the user authentication between Logger and ArcSight Console highly secure. For OTP option to work, Logger must be running 5.1 or later, and the ArcSight Console must be running ESM 5.0 SP1 Patch 2 or later, as described in [“Prerequisites” on page 652](#).

By default, a Logger search from the ArcSight Console uses the OTP method to authenticate. However, if Logger or ArcSight Console is not running a release that supports the OTP option, an error message is displayed and basic authentication is used.

Prerequisites

The following table lists the minimum and recommended versions that Logger and ArcSight Console must be running.

Option	Requirement
Recommended	<p>Logger 5.1 (appliance and software)</p> <p>ESM 5.0 SP1 Patch 2</p> <p>By default, OTP authentication is used. These release versions must be installed for OTP authentication to work.</p>
Minimum	<p>Logger 4.0 or later (on appliance only)</p> <p>ESM 4.5 SP1 Patch 2 or later</p> <p>Basic authentication is used.</p>

Setup and Configuration

ESM

Follow these instructions to set up and configure ArcSight Manager to run integrated search operations:

- 1 Ensure that the ArcSight Manager is running one of the recommended versions. (See ["Prerequisites" on page 652.](#))
- 2 Follow instructions in the ArcSight ESM User's Guide to set up ArcSight Console for integrated searches on Logger. When setting up a user for Logger access (as described in the "Set Up Users for Logger Access" section of the User's Guide), specify the following integration parameters, also shown in the figure below.

Parameter	Type	Value	Targets
OTPPassword	Password	••••••••	Logger Appliance 1
LoggerHost	Text	192.168.36.29	Logger Appliance 1
OTPUser	Text	logger_user	Logger Appliance 1
LoggerPort	Text	443	Logger Appliance 1
LoggerUser	Text	logger_user	Logger Appliance 1
LoggerPassword	Password	••••••••	Logger Appliance 1

Parameter	Description
LoggerUser	<p>The Logger user account to use when accessing a Logger target.</p> <p>For software Logger, this parameter is not applicable. Only OTP method is supported.</p>
LoggerPassword	<p>The password for that Logger account.</p> <p>For software Logger, this parameter is not applicable. Only OTP method is supported.</p>
LoggerHost	The IP address of the Logger host.
OTPUser	The Logger user account to use with the OTP authentication. This account must exist on the Logger.
OTPPassword	The password to use for the OTPUser specified above.
LoggerPort	<p>For OTP, you must specify the port number for the Logger.</p> <p>For a Logger appliance, the port number is 443.</p> <p>For software Logger, specify the port you configured on it during installation.</p>

The ArcSight ESM User's Guide is available from the Protect 724 Community at <https://protect724.arcsight.com>.

Logger

Make sure:

- 1 Your Logger appliance is running a version listed in [“Prerequisites” on page 652](#).
- 2 A Logger user name that you specified when creating an integration parameter on ArcSight Console ([Step 2 of ESM in “Setup and Configuration” on page 653](#)) exists on the Logger.

Supported Search Options

You can select from the following search options when running a Logger search (not Quick Search) from the ArcSight Console:

- By Destination
- By Event Name
- By Source
- By Source and Destination
- By User
- By Vendor and Product

Additionally, if multiple Loggers are configured on your ArcSight Console, you can select the one on which the integrated search should be run.

Guidelines

Make sure you are aware of the following guidelines about Logger Search when it is run from ArcSight Console:

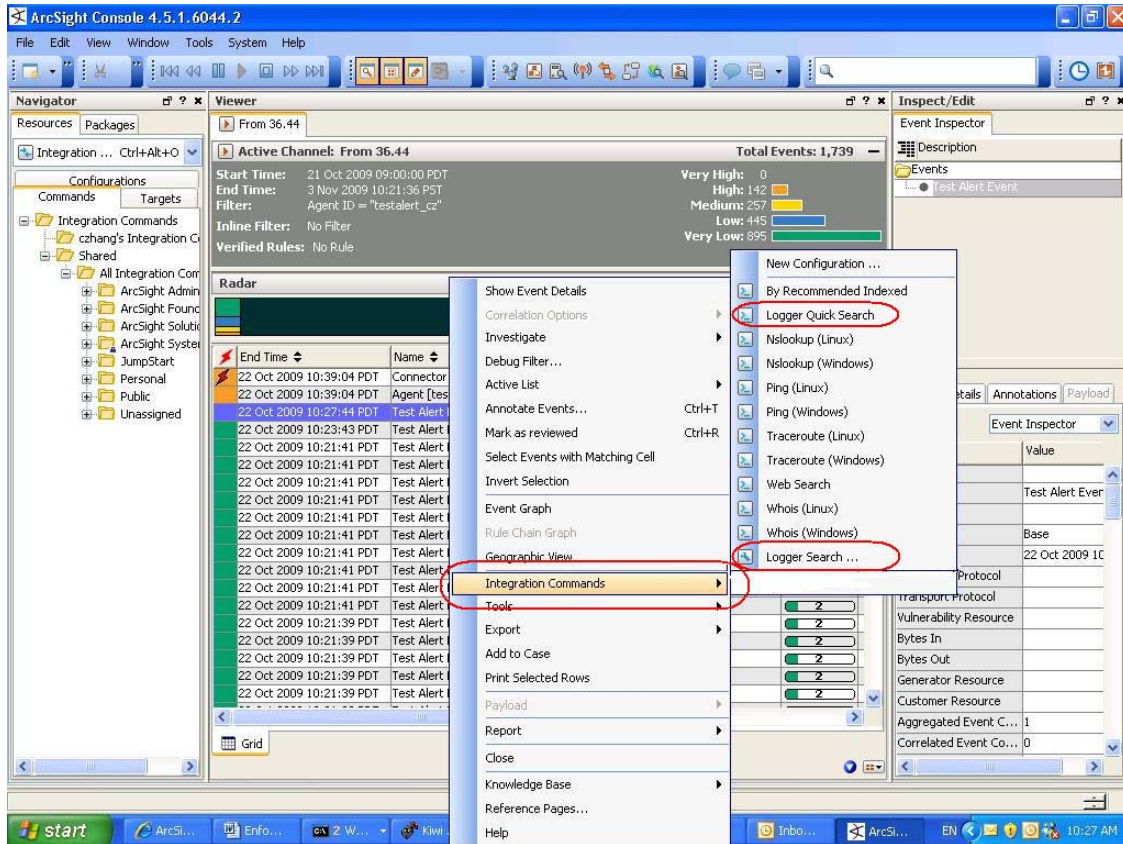
- A field-based search query is used to perform search on the Logger.
- A search operation only from an active channel of an ArcSight Console is supported; search operation from other ESM resources is not supported.
- Multiple search options (see [“Supported Search Options” on page 654](#)) cannot be specified for one search operation. That is, you cannot select by Event Name and By Destination for one search operation.
- You can run the search operation on only one Logger at a time. That is, you cannot select multiple Loggers from the menu list on ArcSight Console.

Additionally, even if a Logger peers with other Loggers, integrated search runs only on the local Logger that you select from the menu list on ArcSight Console.

- The One-Time Password (OTP) authentication is available for use only when Logger is running 5.1 or later and ArcSight Console is running 5.0 SP1 Patch 2 or later. If OTP cannot be used, the searches run from the ArcSight Console display a message that a single-use session token could not be negotiated thus regular authentication will be used. Click OK in that message window so the LoggerUser and LoggerPassword is used to authenticate.

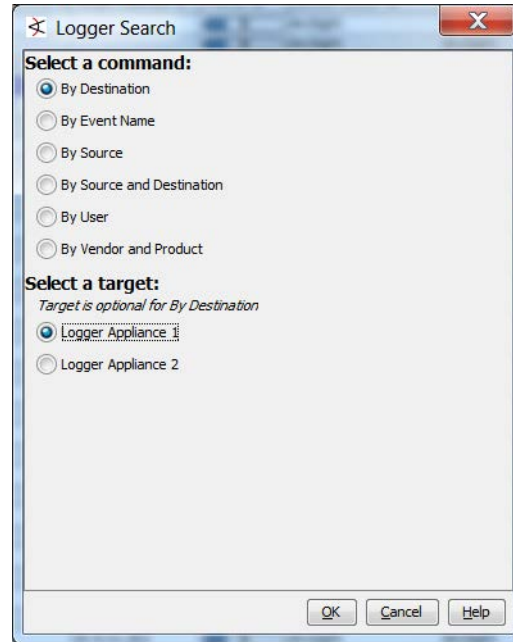
Searching on Logger From ArcSight Console

- To run a Quick Search on Logger (as described in [“Understanding the Integrated Search Functionality”](#) on page 651):
 - a Right click on the event field in an active channel of the ArcSight Console.
 - b From the menu list, select **Integration Commands > Logger Quick Search**, as shown in the following figure.



OR

- To run a regular **Search** (in which you specify search options):
 - a Right click on any field of an event in an active channel of the ArcSight Console.
 - b From the menu list, select **Integration Commands > Logger Search > Select Search Options**, as shown in the following figure.



- c** Click **OK** to run the search or **Cancel** to quit.
- d** If Logger or ArcSight Console is not running a release that supports the OTP option, an error message is displayed indicating that a single-use session token was not negotiated and basic authentication will be used instead.

When such a message is displayed, click OK to proceed.

The search results are displayed in the ArcSight Console Web Viewer.

Index

Numerics

- 1 attempts to gain access through existing accounts 619
- 2 failed file or resource access attempts 620
- 3 unauthorized changes to users groups and services 621
- 32-bit compatibility libraries not found 42
- 4 systems most vulnerable to attack 622
- 5 suspicious or unauthorized network traffic patterns 623

A

- Access Events by Resource query
 - parameters
 - categoryObjectParameter 613
- Access Events by Resource report 613
- Account Modifications report 621
- Accounts Created by User Account report 604
- Accounts Deleted by Host report 604
- Accounts Deleted by User Account report 605
- accounts, user. See users.
- Acronis True Image Server 647
- adhoc reports
 - timeout 238
- advanced mode, packaging connectors 500
- Alert Counts by Device report 600
- Alert Counts by Port report 600
- Alert Counts by Severity report 600
- Alert Counts byType report 600
- Alert Counts per Hour report 600
- alerts
 - about 296
 - adding 299
 - device event class id 566
 - disabling 301, 306
 - enabling 301, 306
 - real time 296
 - remove 301
 - saved search 296
- Alerts from IDS report 624
- anti-virus 594
- anti-virus category 594
- Anti-Virus Updates-All-Failed report 605
- Anti-Virus Updates-All-Summary report 605
- Apache URL Access error log 33
- appliance
 - time, date, and time zone 361
- appliances
 - reimage LX400 and older appliance models 647
 - reimage LX500 645
- ArcExchange 499
- archives

- device event class id 574
- scheduled 254
- ArcSight ESM 22, 27, 54, 56, 285, 295
- ArcSight ESM forwarder 285
- ArcSight Logger Streaming SmartConnector 285
- ArcSight Manager 54, 56, 285, 293
- Asset Startup and Shutdown Event Log report 605
- Attack Events By Destination query
 - parameters
 - destinationAddress 614
 - zones 614
- Attack Events by Destination report 614
- attackers 610
- attackers category 610
- attempts to gain access through existing accounts
 - category 619
- audit events 559
- audit forwarding 559
- audit log 33
- audit.log 33
- AUP file for content update 515
- AUP upgrade process 514
- authentication
 - CAC 396, 432
 - client certificate 396, 432
 - LDAP 396, 432
 - RADIUS 396, 432
- Authentication Errors report 603

B

- backup 330
- Bandwidth Usage by Hour query
 - parameters
 - dmBandwidthParameter 595
- Bandwidth Usage by Hour report 595
- Bandwidth Usage by Protocol query
 - parameters
 - dmBandwidthParameter 595
- Bandwidth Usage by Protocol report 595
- basic mode, packaging connectors 500
- batching 637
- Blocked Ports prompt
 - commonlyBlockedPorts 627
- Bottom Destinations report 614
- Bottom Sources report 610
- Bottom Targets report 615
- browser requirements 57
- bulk copy (see cloning) 526
- By User Account - Accounts Created report 596

C

- CA certificate
 - applying on container 463, 467
 - demo 466
 - invalid errors 469
 - managing 463
 - removing from container 464
 - viewing list 468
 - CAC support 380, 386, 417, 422
 - CACERTS for ESM Destination 294
 - canonical equality check 60
 - cases 504
 - case-sensitive search 60
 - categories 147
 - 1 attempts to gain access through existing accounts 619
 - 2 failed file or resource access attempts 620
 - 3 unauthorized changes to users groups and services 621
 - 4 systems most vulnerable to attack 622
 - 5 suspicious or unauthorized network traffic patterns 623
 - adding a new category 151
 - anti-virus 594
 - attackers 610
 - configuration monitoring 604
 - crossdevice 595
 - database 599
 - deleting an existing one 152
 - device monitoring 594
 - firewall 599
 - foundation 604
 - identity management 601
 - IDS-IPS 600
 - intrusion monitoring 607
 - netflow monitoring 616
 - network 601
 - network monitoring 617
 - operating system 602
 - placing a system defined object in 153
 - resource access 612
 - SANS top 5 619
 - targets 614
 - user tracking 616
 - VPN 603
- Categories tab 642
 - Category Device Group prompt
 - deviceGroupParameter 628
 - categoryObjectParameter parameter 613, 620, 621, 627
 - CEF 129
 - event filters 129
 - events 291
 - TCP receiver 261
 - UDP receiver 261
 - certificate revocation list 387, 424
 - Certificate Signing Request (CSR) 383, 419
 - certificates
 - device event class id 568
 - changing container credentials 461
 - character encoding 267, 268, 270, 272, 273
 - CIFS, configuring 475
 - CLI 410
 - list of commands 410
 - client certificate authentication 396, 432
 - cloning connectors 526
 - Comma Separated Values file, uploading 450
 - command
 - diag sslcert 411
 - exit 410
 - halt 410
 - help 410
 - reboot 410
 - reset authentication 410
 - reset sslcert 411
 - restart process 411
 - restart sslcert 411
 - set date 410
 - set defaultgw 410
 - set dns 411
 - set hostname 411
 - set ip 411
 - set ntp 411
 - set password 411
 - show admin 410
 - show config 410
 - show date 410
 - show defaultgw 410
 - show dns 410
 - show hostname 411
 - show ip 411
 - show ntp 411
 - show sslcert 411
 - show status 411
 - start process 411
 - status process 411
 - stop process 411
 - command line interface (CLI) 410
 - command-line installation 43
 - Common Access Card (CAC) 386, 396, 422, 432
 - common event format (CEF) 129
 - commonlyBlockedPorts parameter 608, 627
 - Configuration - Configuration Changes (Unified) filter 631
 - Configuration - System Configuration Changes (CEF format) filter 631
 - configuration backup 330
 - creating 332
 - device event class id 569
 - guidelines 332
 - restoring 333
 - settings 332
 - Configuration Changes by Type query
 - parameters
 - dmConfigurationParameter 596
 - Configuration Changes by Type report 596
 - Configuration Changes by User query
 - parameters
 - dmConfigurationParameter 596
 - Configuration Changes by User report 596
 - configuration monitoring 604
 - configuration monitoring category 604
 - configuration monitoring, reports for 148
 - Configuration tab 247
 - connect
 - web browsers 57
 - Connection Counts by User report 601, 603
 - Connections Accepted by Address report 603
 - Connections Denied by Address report 603

- Connections Denied by Hour report 603
- Connector Appliance
 - remote upgrade 514
- connector forwarder 285
- connectors supported 474
- constraints, search 78
- containers
 - adding 459
 - changing credentials 461
 - definition 457
 - deleting 459
 - editing 459
 - running commands 470
 - updating properties 460
 - upgrading 470
 - viewing all 458
 - viewing logs 471
- content AUP 515
- content export 353
 - guidelines 355
- content import 353
 - guidelines 354
- copying (see cloning) 526
- crossdevice 595
- crossdevice category 595
- CSR
 - generating a certificate signing request 383, 420
- CSV files
 - information 450
 - uploading 450
- custom connector 499
- Customer URI 639

D

- Daily Bandwidth Usage report 617
- daily data
 - exceeding license limit 41
 - Logger Appliance 34
 - Software Logger 40
- dashboard reports 153, 162
 - preference for display 171
- dashboards
 - device event class id 575
- data limit
 - Logger Appliance 34
 - Software Logger 40
 - violation 41
- database 599
- database category 599
- Database Errors and Warnings report 599
- date/time format 274
- default storage group 21, 52
- demo certificate 466
- Denied Connections by Address report 600
- Denied Connections by Port report 600
- Denied Connections per Hour report 600
- Destination Counts by Device Severity query
 - parameters
 - deviceSeverityParameter 615
- Destination Counts by Device Severity report 615
- Destination Counts by Event Name query
 - parameters
 - eventNameParameter 615
- Destination Counts by Event Name report 615

- Destination IP Address prompt
 - destinationAddress 628
- Destination Port prompt
 - destinationPort 628
- destinationAddress parameter 610, 614, 628
- destinationPort parameter 628
- Device Configuration Changes query
 - parameters
 - deviceGroupParameter 605
- Device Configuration Changes report 605
- Device Configuration Events query
 - parameters
 - deviceGroupParameter 606
- Device Configuration Events report 606
- Device Critical Events report 602
- Device Errors report 602
- device event class ids 560
 - alerts 566
 - archives 574
 - certificates 568
 - configuration backup 569
 - dashboards 575
 - devices 576
 - ESM destinations 569
 - filters 576
 - forwarders 570
 - groups 576
 - maintenance mode 580
 - parsers 577
 - platform 560
 - receivers 572
 - saved searches 578
 - search 580
 - SNMP destinations 573
 - source types 578
 - storage groups 579
 - storage rules 579
 - storage volume 579
 - syslog destinations 573
- Device Events report 602
- device groups 249
 - creating 249
 - deleting 250
 - editing 250
 - maximum number 249
- Device Interface Down Notifications report 602, 607
- Device Interface Status Messages report 602
- Device Misconfigurations query
 - parameters
 - deviceGroupParameter 606
- Device Misconfigurations report 606
- device monitoring 594
- device monitoring category 594
- Device Product prompt
 - deviceProduct 628
- Device Severity prompt
 - deviceSeverityParameter 628
- Device SNMP Authentication Failures report 602
- Device Type prompt
 - dmBandwidthParameter 629
 - dmConfigurationParameter 629
 - dmLoginParameter 629
- Device Vendor prompt
 - deviceVendor 629
- deviceGroupParameter parameter 605, 606, 628

deviceProduct parameter 606, 607, 628

devices 247

- delete 248

- device event class id 576

- edit 248

- maximum number 248

- pre-defining 248

deviceSeverityParameter parameter 611, 615, 628

deviceVendor parameter 606, 607, 629

dmBandwidthParameter parameter 595, 599, 629

dmConfigurationParameter parameter 596, 629

dmLoginParameter parameter 596, 597, 598, 629

dynamic search 87

E

e-mailing reports 182

encoding 267, 268, 270, 272, 273

Errors Detected in Anti-Virus Deployment report 594

ESM (ArcSight Enterprise Security Manager) 22, 27, 54, 56, 285

ESM Destination 291

ESM Destinations 291

- creating 292

- deleting 293

- updating CACERTS 294

ESM destinations

- device event class id 569

eth0 495

event archival, scheduled 254

event archives 250, 253

- adding 253

- deleting 254

- loading 256

- settings 255

- unloading 256

event input 260, 278

- folder follower receivers 261

- receivers 260

- source type 274

Event Name prompt

- eventNameParameter 630

event output 284

- ArcSight ESM 295

- ESM destinations 291

- forwarders 285

event storage, remote 260

eventNameParameter parameter 615, 630

events

- searching 76

Events - CEF filter 631

Events - Event Counts by Destination filter 631

Events - Event Counts by Source filter 631

Events - High and Very High Severity CEF Events filter 631

Events - High and Very High Severity Events (Unified) filter 631

explorers 139

- Category Explorer 140

- Favorites Explorer 144

- icons 144

- Parameter Explorer 143

- Query Explorer 142

- Report Explorer 141

export

- alerts 355

- dashboards 355

- filters 355

- Logger content 355

- parsers 355

- saved searches 355

- search results 121

- source types 355

exporting Logger content 355

exporting remote management configuration 449

extract parser 280

F

factory reset

- LX400 and older appliances 647

- LX500 appliances 645

Failed Anti-Virus Updates report 594

failed file or resource access attempts category 620

Failed Login Attempts query

- parameters

 - dmLoginParameter 596

Failed Login Attempts report 596

Failed Logins by Destination Address query

- parameters

 - dmLoginParameter 596

Failed Logins by Destination Address report 596

Failed Logins by Source Address query

- parameters

 - dmLoginParameter 597

Failed Logins by Source Address report 597

Failed Logins by User query

- parameters

 - dmLoginParameter 597

Failed Logins by User report 597

Failed Res Access Events query

- parameters

 - categoryObjectParameter 621

Failed Resource Access by Users report 620

Failed Resource Access Events report 621

Failed Resource Access query

- parameters

 - categoryObjectParameter 620

field operators 80

field query

- indexing fields 122

field set, search 78

fields, indexing 122

file receivers 261, 262

- multi-line 262

file transfer receiver 261, 262

filtering information on UI page 447

filters 312, 631

- Configuration - Configuration Changes (Unified) 631

- Configuration - System Configuration Changes (CEF format) 631

- copying 313

- creating 312

- deleting 313

- device event class id 576

- editing 313

- Events - CEF 631

- Events - Event Counts by Destination 631

- Events - Event Counts by Source 631

- Events - High and Very High Severity CEF Events 631
- Events - High and Very High Severity Events (Unified) 631
- Firewall - Deny 632
- Firewall - Drop 632
- Firewall - Permit 632
- Intrusion - Malicious Code (CEF format) 632
- Intrusion - Malicious Code (Unified) 632
- Logins - All Logins (CEF format) 632
- Logins - All Logins (Non-CEF format) 632
- Logins - All Logins (Unified) 632
- Logins - Failed Logins 632
- Logins - Successful Logins (CEF format) 632
- Logins - Successful Logins (Non-CEF format) 632
- Logins - Successful Logins (Unified) 632
- Logins - Unsuccessful Logins (CEF format) 632
- Logins - Unsuccessful Logins (Non-CEF format) 632
- Logins - Unsuccessful Logins (Unified) 632
- Network - DHCP Lease Events 632
- Network - Port Links Up and Down 632
- Network - Protocol Links Up and Down 632
- report category 239
- search 78
- system 128
- SystemAlert - Bad Fan (CEF format) 633
- SystemAlert - CPU Utilization Above 90% (CEF format) 632
- SystemAlert - CPU Utilization Above 90% (Unified) 633
- SystemAlert - CPU Utilization Above 95% (CEF format) 633
- SystemAlert - CPU Utilization Above 95% (Unified) 633
- SystemAlert - Device Configuration Changes (CEF format) 633
- SystemAlert - Device Configuration Changes (Unified) 633
- SystemAlert - Disk Failure (CEF format) 633
- SystemAlert - Disk Failure (Unified) 633
- SystemAlert - Filter Configuration Changes (CEF format) 633
- SystemAlert - Filter Configuration Changes (Unified) 633
- SystemAlert - High CPU Temperature (CEF format) 633
- SystemAlert - High CPU Temperature (Unified) 633
- SystemAlert - Power Supply Failure (CEF format) 633
- SystemAlert - Power Supply Failure (Unified) 633
- SystemAlert - RAID Controller Issue (CEF format) 633
- SystemAlert - RAID Controller Issue (Unified) 634
- SystemAlert - RAID Status Battery Failure (CEF format) 633
- SystemAlert - RAID Status Battery Failure (Unified) 633
- SystemAlert - Root Partition Free Space Below 10% (CEF format) 634
- SystemAlert - Root Partition Free Space Below 10% (Unified) 634
- SystemAlert - Root Partition Free Space Below 5% (CEF format) 634
- SystemAlert - Root Partition Free Space Below 5% (Unified) 634
- SystemAlert - Storage Configuration Changes (CEF format) 634
- SystemAlert - Storage Configuration Changes (Unified) 634
- SystemAlert - Storage Group Usage Above 90% (CEF format) 634
- SystemAlert - Storage Group Usage Above 90% (Unified) 634
- SystemAlert - Storage Group Usage Above 95% (CEF format) 634
- SystemAlert - Storage Group Usage Above 95% (Unified) 634
- SystemAlert - Zero Events Incoming (CEF format) 634
- SystemAlert - Zero Events Incoming (Unified) 634
- SystemAlert - Zero Events Outgoing (CEF format) 634
- SystemAlert - Zero Events Outgoing (Unified) 635
- SystemStatus - CPU Utilization by Connector Host 635
- SystemStatus - Disk Utilization by Connector Host 635
- SystemStatus - Memory Utilization by Connector Host 635
- Unix - CRON related events 635
- Unix - IO Errors and Warnings 635
- Unix - PAM and Sudo Messages 635
- Unix - Password Changes 635
- Unix - SAMBA Events 635
- Unix - SSH Authentications 635
- Unix - User and Group Additions 635
- Unix - User and Group Deletions 635
- Windows - Account Added to Global Group 635
- Windows - Account Added to Global Group (CEF) 635
- Windows - Audit Policy Change 635
- Windows - Audit Policy Change (CEF) 635
- Windows - Change Password Attempt 635
- Windows - Change Password Attempt (CEF) 635
- Windows - Global Group Created 635
- Windows - Global Group Created (CEF) 635
- Windows - Logon Bad User Name or Password 636
- Windows - Logon Bad User Name or Password (CEF) 636
- Windows - Logon Local User 636
- Windows - Logon Local User (CEF) 636
- Windows - Logon Remote User 636
- Windows - Logon Remote User (CEF) 636
- Windows - Logon Unexpected Failure 636
- Windows - Logon Unexpected Failure (CEF) 636
- Windows - New Process Creation 636
- Windows - New Process Creation (CEF) 636
- Windows - Pre-Authentication Failure 636
- Windows - Pre-Authentication Failure (CEF) 636
- Windows - Special Privileges Assigned to New Logon 636
- Windows - Special Privileges Assigned to New Logon (CEF) 636
- Windows - User Account Changed 636
- Windows - User Account Changed (CEF) 636
- Windows - User Account Password Set 636
- Windows - User Account Password Set (CEF) 636
- Windows - Windows Events (CEF) 636
- finding events 76
- FIPS 140-2

- enabling on Connector Appliance 388, 424
- enabling on container 461
- Firefox (web browser) 57
- firewall 599
- Firewall - Deny filter 632
- Firewall - Drop filter 632
- Firewall - Permit filter 632
- firewall category 599
- Firewall Traffic by Service report 607
- folder follower receivers 33, 261, 262
 - /opt/local/apache/logs/http_error_log 33
 - /userdata/logs/apache/http_error_log 33
 - /var/log/audit/audit.log 33
 - /var/log/messages 33
- Logger Appliance 33
- forgot password 396, 432
- forwarder types
 - ArcSight ESM 285
 - connector 285
 - TCP 285
 - UDP 285
- forwarders 285
 - creating 286
 - deleting 290
 - device event class id 570
 - editing 289
- forwarding file events to ESM 295
- foundation 604
- foundation category 604
- function tabs 59

G

- gauge range 60
- gauges 59
- glibc-2.12-1.25.el6.i686 42
- groups, device event class ID 576

H

- health, system 407, 443
- help 59, 60
- hosts
 - adding 452
 - definition 451
 - deleting 456
 - editing 456
 - moving to different location 456
 - scanned 453
 - scanning 454
 - software-type 451
 - upgrading remotely 457
 - viewing all 452
- Hourly Bandwidth Usage report 617
- http_error_log 33

I

- i18n options 60
- identity management 601
- identity management category 601
- IDS Signature Destinations report 624
- IDS Signature Sources report 624
- IDS-IPS 600
- IDS-IPS category 600

- import
 - alerts 354
 - dashboards 354
 - filters 354
 - Logger content 354
 - parsers 354
 - remote management configuration 449
 - saved searches 354
 - source types 354
- increasing search speed 101
- indexing fields 122
- install software Logger 43
- installation wizard 42
- installing Logger on VMWare VM 29
- internal storage group 21, 257
- internationalization options 60
- Internet Explorer (web browser) 57
- Intrusion - Malicious Code (CEF format) filter 632
- Intrusion - Malicious Code (Unified) filter 632
- intrusion monitoring 607
 - reports for 148
- intrusion monitoring category 607
- invalid certificate errors 469
- IP Address prompt
 - IPAddress 627
- IPAddress parameter 611, 616, 623, 627

L

- launch Logger 48
- LDAP 477
- LDAP authentication 396, 432
- Least Common Accessed Ports report 613
- Least Common Events report 607
- license
 - Logger Appliance 34
 - Software Logger 40
- license information 34, 40, 41
- Localhost 451
- localization options 60
- locations
 - adding 448
 - definition 447
 - deleting 451
 - editing 451
 - viewing all 448
- log in 34, 49, 57, 58
- Logfu utility 472
- Logger Appliance
 - initialization 29
- Logger content
 - exporting 355
 - importing 354
- Logger on VMWare
 - installation 29
- loggerd 48
 - quit 48, 49
 - restart 49
 - start 48, 49
 - status 49
 - stop 48, 49
- login banner 402, 438
- Login Errors by User report 603
- Login Event Audit query
 - parameters

- dmLoginParameter 597
- Login Event Audit report 597
- login screen 49, 58
- Logins - All Logins (CEF format) filter 632
- Logins - All Logins (Non-CEF format) filter 632
- Logins - All Logins (Unified) filter 632
- Logins - Failed Logins filter 632
- Logins - Successful Logins (CEF format) filter 632
- Logins - Successful Logins (Non-CEF format) filter 632
- Logins - Successful Logins (Unified) filter 632
- Logins - Unsuccessful Logins (CEF format) filter 632
- Logins - Unsuccessful Logins (Non-CEF format) filter 632
- Logins - Unsuccessful Logins (Unified) filter 632
- logout 59, 61
 - automatic 61
- logs
 - deleting container 472
 - internal 352
 - internal, retrieving 352
 - repository 511
 - uploading to repository 511
 - viewing container 471
- LX400 and older appliances, reimaging 647
- LX500 appliances, reimaging 645

M

- maintenance mode 333
 - device event class id 580
- monitor a directory 262
- Monitor tab 70
- Most Common Events by Severity report 608
- Most Common Events report 608
- multi-line receivers 262

N

- navigation 59
- needle-in-a-haystack searches 83
- netflow monitoring 616
- netflow monitoring category 616
- network 601
- Network - DHCP Lease Events filter 632
- Network - Port Links Up and Down filter 632
- Network - Protocol Links Up and Down filter 632
- network category 601
- network interfaces 495
- network monitoring 617
- network monitoring category 617
- NFS, configuring 475
- nss-softokn-freebl-3.12.9-3.el6.i686 42
- Number of Failed Logins report 619

O

- online help 60
- operating system 602
- operating system category 602
- operators
 - field 80
- options 59, 60
- original factory settings
 - restore LX400 and older appliance models 647
 - restore LX500 645

P

- packaging connectors
 - advanced mode 500
 - basic mode 500
- parameter value groups, in reports 229
- parameters 626
 - categoryObjectParameter 613, 620, 621, 627
 - commonlyBlockedPorts 608, 627
 - destinationAddress 610, 614, 628
 - destinationPort 628
 - deviceGroupParameter 605, 606, 628
 - deviceProduct 606, 607, 628
 - deviceSeverityParameter 611, 615, 628
 - deviceVendor 606, 607, 629
 - dmBandwidthParameter 595, 599, 629
 - dmConfigurationParameter 596, 629
 - dmLoginParameter 596, 597, 598, 629
 - eventNameParameter 615, 630
 - in report queries 223
 - IPAddress 611, 616, 623, 627
 - quick run report 175
 - resourceTypeParameter 630
 - run reports 178
 - webPorts 609, 614, 630
 - zoneParameter 611, 612, 616, 623, 630
 - zones 610, 614, 631
- parse command 279
- parser override 499
- parsers 278
 - device event class id 577
 - extract 280
 - REX 280
 - using with source types 279
- password
 - changing 407, 442
 - reset 396, 432
- Password Change query
 - parameters
 - dmLoginParameter 597
- Password Changes report 597, 606, 622
- passwords
 - changing 407, 442
- peer Loggers 326
 - adding 328
 - authorizing 330
 - deleting 329
 - searching 108
- peer relationships 326
- peers
 - device event class id 577
- platform
 - device event class id 560
- pre-configured folder follower receivers 33
- predefined filters 128, 129
- Probes on Blocked Ports by Source query
 - parameters
 - commonlyBlockedPorts 608
- Probes on Blocked Ports by Source report 608
- prompts
 - Blocked Ports 627
 - Category Device Group 628
 - Destination IP Address 628
 - Destination Port 628
 - Device Product 628

- Device Severity 628
- Device Type 629
- Device Vendor 629
- Event Name 630
- IP Address 627
- Resource Type 627, 630
- Web Ports 630
- Zone 630, 631
- Protect 724 499

Q

queries

- 1 - Attempts to Gain Access through Existing Accounts 619
- 2 - Failed File or Resource Access Attempts 620
- 3 - Unauthorized Changes to Users Groups and Services 621
- 4 - Systems Most Vulnerable to Attack 623
- 5 - Suspicious or Unauthorized Network Traffic Patterns 624

- Anti-Virus 594
- Attackers 610
- Configuration Monitoring 604
- CrossDevice 595
- Database 599
- Firewall 600
- Identity Management 601
- IDS-IPS 600
- Intrusion Monitoring 607
- NetFlow Monitoring 617
- Network 602
- Network Monitoring 617
- Operating System 603
- Resource Access 613
- Targets 614
- VPN 603

query

- controls 59
- events 76
- use in reports 201

R

- RADIUS authentication 396, 432
- RAID controller status 365, 415
- range, gauge 60
- rare field values, searching for 83
- real time alerts 296
- rebooting 358
- receiver types
 - CEF TCP 261
 - CEF UDP 261
 - file receiver 261, 262
 - file transfer 261
 - file transfer receiver 262
 - folder follower 261
 - folder follower receiver 262
 - multi-line 262
 - SmartMessage 261
 - TCP 261
 - UDP 261
- receivers 260, 264
 - creating 264
 - deleting 266

- device event class id 572
- disabling 265
- editing 265
- enabling 265
- folder follower 262
- refreshing UI screen 447
- regular expressions (regex)
 - predefined 129
- reimage
 - LX400 and older appliance models 647
 - LX500 appliances 645
- Remote Authentication Dial-In User Service (RADIUS) 396, 432
- remote event storage 260
- remote file system mount
 - adding 376
 - editing 377
- remote management configuration 449
 - exporting 449
 - importing 449
- remote upgrade 514
- report administration 238
- report category filters 239
- report package, deploying 236
- report server configuration 238
- reports 316, 593
 - 1 - Attempts to Gain Access through Existing Accounts 619
 - 2 - Failed File or Resource Access Attempts 620
 - 3 - Unauthorized Changes to Users Groups and Services 621
 - 4 - Systems Most Vulnerable to Attack 623
 - 5 - Suspicious or Unauthorized Network Traffic Patterns 624
- Access Events by Resource 613
- access rights 198
- Account Modifications 621
- Accounts Created by User Account 604
- Accounts Deleted by Host 604
- Accounts Deleted by User Account 605
- administration 238
- Alert Counts by Device 600
- Alert Counts by Port 600
- Alert Counts by Severity 600
- Alert Counts by Type 600
- Alert Counts per Hour 600
- Alerts from IDS 624
- Anti-Virus 594
- Anti-Virus Updates-All-Failed 605
- Anti-Virus Updates-All-Summary 605
- Asset Startup and Shutdown Event Log 605
- Attack Events by Destination 614
- Attackers 610
- attackers 619
- Authentication Errors 603
- Bandwidth Usage by Hour 595
- Bandwidth Usage by Protocol 595
- Bottom Destinations 614
- Bottom Sources 610
- Bottom Targets 615
- By User Account - Accounts Created 596
- client timeout 238
- Configuration Changes by Type 596
- Configuration Changes by User 596
- Configuration Monitoring 604

- configuration monitoring 148
- Connection Counts by User 601, 603
- Connections Accepted by Address 603
- Connections Denied by Address 603
- Connections Denied by Hour 603
- creating new 186
- CrossDevice 595
- Daily Bandwidth Usage 617
- dashboard 153, 162
- Database 599
- database connection timeout 238
- Database Errors and Warnings 599
- delivery options 181
- Denied Connections by Address 600
- Denied Connections by Port 600
- Denied Connections per Hour 600
- designing 185
- Destination Counts by Device Severity 615
- Destination Counts by Event Name 615
- Device Configuration Changes 605
- Device Configuration Events 606
- Device Critical Events 602
- Device Errors 602
- Device Events 602
- Device Interface Down Notifications 602, 607
- Device Interface Status Messages 602
- Device Misconfigurations 606
- Device SNMP Authentication Failures 602
- editing 197
- e-mailing 182
- Errors Detected in Anti-Virus Deployment 594
- exporting 184
- Failed Anti-Virus Updates 594
- Failed Login Attempts 596
- Failed Logins by Destination Address 596
- Failed Logins by Source Address 597
- Failed Logins by User 597
- Failed Resource Access by Users 620
- Failed Resource Access Events 621
- file formats 179
- Firewall 600
- Firewall Traffic by Service 607
- Hourly Bandwidth Usage 617
- Identity Management 601
- IDS Signature Destinations 624
- IDS Signature Sources 624
- IDS-IPS 600
- Intrusion Monitoring 607
- intrusion monitoring 148
- iPackager 240
- Least Common Accessed Ports 613
- Least Common Events 607
- Login Errors by User 603
- Login Event Audit 597
- Most Common Events 608
- Most Common Events by Severity 608
- navigating to 137
- NetFlow Monitoring 617
- Network 602
- Network Monitoring 617
- Number of Failed Logins 619
- Operating System 603
- parameter value groups 229
- Password Changes 597, 606, 622
- Probes on Blocked Ports by Source 608
- publishing 180
- query parameters 223
- quick run parameters 175
- recently run 173
- remove scheduled 234
- Resource Access 613
- Resource Access by Users - Failures 613
- Resource Access by Users - Successes-Attempts 613
- run parameters 178
- running 172
- SANS Top 5 149
- saving 184
- scheduling 233
- SecurityDashBoardRpt 608
- SecurityDBReport 609
- solution add-ons 150
- Source Counts by Destination 610
- Source Counts by Destination Port 610
- Source Counts by Device 611
- Source Counts by Device Severity 611
- Source Counts by Source Port 611
- Source Port Counts 611
- Successful Logins by Destination Address 598
- Successful Logins by Source Address 598
- Successful Logins by User 598
- Target Attack Counts by Severity 615
- Target Counts by Event Name 615
- Target Counts by Severity 615
- Target Counts by Source 615
- Target Counts by Source Port 616
- Target Counts by Target Port 616
- Target Port Counts 616
- Targets 614
- template styles 232
- Top 10 Talkers 611, 624
- Top 10 Types of Traffic 625
- Top Alert Destinations 600
- Top Alert Sources 601
- Top Alerts from IDS 625
- Top Alerts from IDS and IPS 601
- Top Attack Sources 612
- Top Attacker Detail 611
- Top Attacker Details 611
- Top Attacker Ports 612
- Top Attackers 612
- Top Bandwidth Hosts 599
- Top Bandwidth Usage by Destination 617
- Top Bandwidth Usage by Destination Port 617
- Top Bandwidth Usage by Source 617
- Top Destination IPs 625
- Top Destination Ports 616
- Top Destinations Across Firewalls 616
- Top Destinations in IDS Events 616
- Top Hosts by Number of Connections 599
- Top IDS Attack Events 609
- Top IDS Events 609
- Top IDS Signature Destinations 625
- Top IDS Signature Sources 626
- Top Infected Systems 594
- Top Machines Accessing the Web 614
- Top Machines Traversing Firewall 609
- Top Sources Detected by Snort 612
- Top Sources Traversing Firewalls 612
- Top Target IPs 626

- Top Targets 616
- Top Users with Failed Logins 620
- Top VPN Accesses by User 617
- Top VPN Event Destinations 618
- Top VPN Event Sources 618
- Top VPN Events 618
- Top Web Traffic 609
- Traffic Statistics 618
- Update Summary 594
- User Account Creations 622
- User Account Deletions 622
- User Account Modifications 622
- User Administration 603
- viewing published 184
- viewing, editing schedules 233
- Virus Activity by Hour 595
- VPN 603
- VPN Connection Attempts 618
- VPN Connection Failures 619
- Vulnerability Scanner Logs by Host 606, 623
- Vulnerability Scanner Logs by Vulnerability 607, 623
- Windows Events 609
- Worm Infected Systems 601, 609
- repositories, user-defined 518
- repository
 - logs 511
- reset password 396, 432
- resource access 612
- Resource Access by Users - Failures query
 - parameters
 - categoryObjectParameter 613
- Resource Access by Users - Failures report 613
- Resource Access by Users - Success-Attempt query
 - parameters
 - categoryObjectParameter 613
- Resource Access by Users - Successes-Attempts report 613
- resource access category 612
- Resource Type prompt
 - categoryObjectParameter 627
 - resourceTypeParameter 630
- resources parameters 626
- resourceTypeParameter parameter 630
- restart 48
- restarting Software Logger 48
- restore LX400 and older appliance models to original factory settings 647
- restore LX500 to original factory settings 645
- restoring a SAN 377
- retrieve logs 352
- REX parser 280

S

- SAN, restore 377
- SANS top 5 619
- SANS top 5 category 619
- SANS Top 5, reports for 149
- saved
 - filters 126
 - search 126
- saved search alerts 296
- saved search files 320
- saved search job 316

- adding 316
- deleting 319
- editing 319
- saved searches 315
 - adding 315
 - deleting 316
 - device event class id 578
 - editing 316
- scan a host 453, 454
- scheduled event archive 254
- scheduled reports
 - timeout 238
- scheduled tasks 309, 310
 - currently running 310
 - finished 311
- scheduling
 - export of search results 108
 - reports 233
- SCP file receivers 270
- search
 - constraints 78
 - defining queries 78
 - device event class id 580
 - events 76
 - exporting results 121
 - field set 78
 - filters 78, 126
 - increasing speed of 83
 - peer Loggers 108
 - results, scheduling export of 108
 - saved 126
 - system filters 128
 - time range 78
- search group filters 314
 - associating with user group 314
 - report category filters 239
- search operators
 - cef (deprecated) 529
 - chart 530
 - dedup 536
 - eval 536
 - extract 537
 - fields 539
 - head 539
 - keys 540
 - parse 541
 - rare 542
 - regex 542
 - rename 543
 - replace 544
 - rex 545
 - sort 547
 - tail 548
 - top 548
 - transaction 549
 - where 551
- Search Results tab 109
- search speed, increasing 83
- searching for rare values in fields 83
- SecurityDashBoardRpt report 608
- SecurityDBReport report 609
- send file events to ESM 295
- severity level 637, 639, 642
- SFTP file receivers 270
- silent installation 43

- SmartConnectors 33, 53, 55, 637, 639
 - batching 637
 - configuring 53
 - defined 474
 - scanner 641
 - zones 639
- SmartMessage receivers 33, 261
 - configuring 53
- SNMP 407, 443, 445, 480, 504
- SNMP destinations
 - device event class id 573
- Software Logger
 - folder follower receivers 33
 - restarting 48
 - starting 48
 - stopping 48
 - time, date, and time zone 361
- software Logger
 - installing 43
 - uninstalling 49
- Software Logger installation 29
- software-type host 451
- solutions
 - reports 150
- Source Counts by Destination Port report 610
- Source Counts By Destination query
 - parameters
 - destinationAddress 610
 - zones 610
- Source Counts by Destination report 610
- Source Counts by Device report 611
- Source Counts by Device Severity query
 - parameters
 - deviceSeverityParameter 611
- Source Counts by Device Severity report 611
- Source Counts by Source Port report 611
- Source Port Counts report 611
- source types
 - device event class id 578
- SSL 381, 383, 417, 419
 - Certificate Signing Request 383, 419
- starting Software Logger 48
- statistics 59
- status
 - 3Ware RAID Controller 365, 415
 - GetStatus command 502, 516, 528
- stopping Software Logger 48
- storage 256
- storage groups 257
 - default 21, 52, 257, 258
 - device event class id 579
 - editing 257
 - internal 21
- storage rules 52, 259
 - adding 259
 - deleting 260
 - device event class id 579
 - editing 260
- storage settings 260
- storage volume 260
 - device event class id 579
- streaming SmartConnector 285
- structured data 22, 52
 - exporting 25
 - searching 86
- Successful Logins by Destination Address query
 - parameters
 - dmLoginParameter 598
- Successful Logins by Destination Address report 598
- Successful Logins by Source Address query
 - parameters
 - dmLoginParameter 598
- Successful Logins by Source Address report 598
- Successful Logins by User query
 - parameters
 - dmLoginParameter 598
- Successful Logins by User report 598
- Super Indexes 125
- super-indexed fields
 - list of 84
 - searching 83
- super-indexed search 83
 - optimizations 84
- supported connectors 474
- suspicious or unauthorized network traffic patterns
 - category 623
- syslog destinations
 - device event class id 573
- system audit log 33
- system definition 446
- system filters 128, 129, 631
- system health, monitoring 407, 443
- system messages log 33
- system reboot 358
- system restore 645, 647
 - LX400 and older appliance models 647
 - LX500 appliances 645
- system restore utility 645
- SystemAlert - Bad Fan (CEF format) filter 633
- SystemAlert - CPU Utilization Above 90% (CEF format) filter 632
- SystemAlert - CPU Utilization Above 90% (Unified) filter 633
- SystemAlert - CPU Utilization Above 95% (CEF format) filter 633
- SystemAlert - CPU Utilization Above 95% (Unified) filter 633
- SystemAlert - Device Configuration Changes (CEF format) filter 633
- SystemAlert - Device Configuration Changes (Unified) filter 633
- SystemAlert - Disk Failure (CEF format) filter 633
- SystemAlert - Disk Failure (Unified) filter 633
- SystemAlert - Filter Configuration Changes (CEF format) filter 633
- SystemAlert - Filter Configuration Changes (Unified) filter 633
- SystemAlert - High CPU Temperature (CEF format) filter 633
- SystemAlert - High CPU Temperature (Unified) filter 633
- SystemAlert - Power Supply Failure (CEF format) filter 633
- SystemAlert - Power Supply Failure (Unified) filter 633
- SystemAlert - RAID Controller Issue (CEF format) filter 633
- SystemAlert - RAID Controller Issue (Unified) filter 634
- SystemAlert - RAID Status Battery Failure (CEF format) filter 633
- SystemAlert - RAID Status Battery Failure (Unified) filter 633

- SystemAlert - Root Partition Free Space Below 10% (CEF format) filter 634
- SystemAlert - Root Partition Free Space Below 10% (Unified) filter 634
- SystemAlert - Root Partition Free Space Below 5% (CEF format) filter 634
- SystemAlert - Root Partition Free Space Below 5% (Unified) filter 634
- SystemAlert - Storage Configuration Changes (CEF format) filter 634
- SystemAlert - Storage Configuration Changes (Unified) filter 634
- SystemAlert - Storage Group Usage Above 90% (CEF format) filter 634
- SystemAlert - Storage Group Usage Above 90% (Unified) filter 634
- SystemAlert - Storage Group Usage Above 95% (CEF format) filter 634
- SystemAlert - Storage Group Usage Above 95% (Unified) filter 634
- SystemAlert - Zero Events Incoming (CEF format) filter 634
- SystemAlert - Zero Events Incoming (Unified) filter 634
- SystemAlert - Zero Events Outgoing (CEF format) filter 634
- SystemAlert - Zero Events Outgoing (Unified) filter 635
- systems most vulnerable to attack category 622
- SystemStatus - CPU Utilization by Connector Host filter 635
- SystemStatus - Disk Utilization by Connector Host filter 635
- SystemStatus - Memory Utilization by Connector Host filter 635

T

- Target Attack Counts by Severity report 615
- Target Counts by Event Name report 615
- Target Counts by Severity report 615
- Target Counts by Source Port report 616
- Target Counts by Source report 615
- Target Counts by Target Port report 616
- Target Port Counts report 616
- targets 614
- targets category 614
- TCP forwarders 285
- TCP receivers 261
 - default port 32
- template styles for reports 232
- time range
 - dynamic 87
 - search 78
- time, date, and time zone
 - configuring on Logger Appliance 361
 - configuring on Software Logger 361
- Top 10 Talkers report 611, 624
- Top 10 Types of Traffic report 625
- Top Alert Destinations report 600
- Top Alert Sources report 601
- Top Alerts from IDS and IPS report 601
- Top Alerts from IDS report 625
- Top Attack Sources report 612
- Top Attacker Detail report 611
- Top Attacker Details query
 - parameters

- IPAddress 611
- zoneParameter 611
- Top Attacker Details report 611
- Top Attacker Ports report 612
- Top Attackers query
 - parameters
 - zoneParameter 612
- Top Attackers report 612
- Top Bandwidth Hosts query
 - parameters
 - dmBandwidthParameter 599
- Top Bandwidth Hosts report 599
- Top Bandwidth Usage by Destination Port report 617
- Top Bandwidth Usage by Destination report 617
- Top Bandwidth Usage by Source report 617
- Top Destination IPs report 625
- Top Destination Ports report 616
- Top Destinations Across Firewalls report 616
- Top Destinations in IDS Events report 616
- Top Hosts by Number of Connections query
 - parameters
 - dmBandwidthParameter 599
- Top Hosts by Number of Connections report 599
- Top IDS Attack Events report 609
- Top IDS Events report 609
- Top IDS Signature Destinations report 625
- Top IDS Signature Sources report 626
- Top Infected Systems report 594
- Top Machines Accessing the Web query
 - parameters
 - webPorts 614
- Top Machines Accessing the Web report 614
- Top Machines Traversing Firewall report 609
- Top Sources Detected by Snort report 612
- Top Sources Traversing Firewalls report 612
- Top Target IPs report 626
- Top Targets query
 - parameters
 - IPAddress 616
 - zoneParameter 616
- Top Targets report 616
- Top Users with Failed Logins report 620
- Top VPN Accesses by User report 617
- Top VPN Event Destinations report 618
- Top VPN Event Sources report 618
- Top VPN Events report 618
- Top Web Traffic query
 - parameters
 - webPorts 609
- Top Web Traffic report 609
- Traffic Statistics report 618
- trial license 41

U

- UDP forwarders 285
- UDP receivers 261
 - default port 32
- unauthorized changes to users groups and services
 - category 621
- Unicode options 60
- uninstalling Logger software 49
- Unix - CRON related events filter 635
- Unix - IO Errors and Warnings filter 635
- Unix - PAM and Sudo Messages filter 635

- Unix - Password Changes filter 635
- Unix - SAMBA Events filter 635
- Unix - SSH Authentications filter 635
- Unix - User and Group Additions filter 635
- Unix - User and Group Deletions filter 635
- unstructured data 22
 - exporting 25
 - searching 86
- Update Summary report 594
- update, content 515
- updating container properties 460
- upgrade
 - Connector Appliance 514
 - host 514
 - remote 514
- US-ASCII encoding 267, 268, 270, 272, 273
- User Account Creations report 622
- User Account Deletions report 622
- User Account Modifications report 622
- User Administration report 603
- user groups
 - associating with search group filters 314
 - creating 406, 441
 - deleting 407, 442
 - editing 406, 442
- user interface 59
 - filtering information to display 447
 - refresh 447
 - Search Results tab 109
- user rights for content export
 - content export
 - user rights 354
- user rights for content import
 - content import
 - user rights 353
- user tracking 616
- user tracking category 616
- user-defined repositories 518
- users
 - changing password 407, 442
 - creating 402, 438
 - deleting 404, 440
 - editing 403, 439
- UTF-8 encoding 267, 268, 270, 272, 273

V

- Virus Activity by Hour report 595
- VPN 603
- VPN category 603
- VPN Connection Attempts report 618
- VPN Connection Failures report 619
- Vulnerability Scanner Logs by Host query
 - parameters
 - deviceProduct 606
 - deviceVendor 606
- Vulnerability Scanner Logs by Host report 606, 623
- Vulnerability Scanner Logs by Vulnerability query
 - parameters

- deviceProduct 607
 - deviceVendor 607
- Vulnerability Scanner Logs by Vulnerability report 607, 623
- Vulnerability Scanner Logs query
 - parameters
 - IPAddress 623
 - zoneParameter 623

W

- web browser requirements 57
- Web Ports prompt
 - webPorts 630
- webPorts parameter 609, 614, 630
- what's new 28
- widgets in report dashboards 168
- Windows - Account Added to Global Group (CEF) filter 635
- Windows - Account Added to Global Group filter 635
- Windows - Audit Policy Change (CEF) filter 635
- Windows - Audit Policy Change filter 635
- Windows - Change Password Attempt (CEF) filter 635
- Windows - Change Password Attempt filter 635
- Windows - Global Group Created (CEF) filter 635
- Windows - Global Group Created filter 635
- Windows - Logon Bad User Name or Password (CEF) filter 636
- Windows - Logon Bad User Name or Password filter 636
- Windows - Logon Local User (CEF) filter 636
- Windows - Logon Local User filter 636
- Windows - Logon Remote User (CEF) filter 636
- Windows - Logon Remote User filter 636
- Windows - Logon Unexpected Failure (CEF) filter 636
- Windows - Logon Unexpected Failure filter 636
- Windows - New Process Creation (CEF) filter 636
- Windows - New Process Creation filter 636
- Windows - Pre-Authentication Failure (CEF) filter 636
- Windows - Pre-Authentication Failure filter 636
- Windows - Special Privileges Assigned to New Logon (CEF) filter 636
- Windows - Special Privileges Assigned to New Logon filter 636
- Windows - User Account Changed (CEF) filter 636
- Windows - User Account Changed filter 636
- Windows - User Account Password Set (CEF) filter 636
- Windows - User Account Password Set filter 636
- Windows - Windows Events (CEF) filter 636
- Windows Events report 609
- Worm Infected Systems report 601, 609

Z

- Zone prompt
 - zoneParameter 630
 - zones 631
- zoneParameter parameter 611, 612, 616, 623, 630
- zones parameter 610, 614, 631

